StarTeam Administrator's Guide

StarTeam®

Borland®
Excellence Endures™

Borland Software Corporation
100 Enterprise Way
Scotts Valley, California 95066-3249
www.borland.com

Borland Software Corporation may have patents and/or pending patent applications covering subject matter in this document. Please refer to the product CD or the About dialog box for the list of applicable patents. The furnishing of this document does not give you any license to these patents.

COPYRIGHT © 1995–2005 Borland Software Corporation. All rights reserved. All Borland brand and product names are trademarks or registered trademarks of Borland Software Corporation in the United States and other countries. All other marks are the property of their respective owners.

For third-party conditions and disclaimers, see the Release Notes on your product CD.

Printed in the U.S.A.

Contents

Chapter 1
Preface **1**
Application Documentation Set 1
Online Documentation 2
Contacting Borland Support 3
Documentation Conventions 3

Chapter 2
Administering the Application **5**
Understanding Server Configurations. 6
Understanding This Guide 6

Chapter 3
Configuring and Running the Server **9**
Planning Server Configurations. 9
Managing Server Configurations with the Server Administration Utility. 11
Starting the Server Administration Utility 11
Understanding the starteam-servers.xml File . . . 12
Creating a Server Configuration 12
Starting a Server Configuration 15
Logging on as a Different User 16
Shutting Down a Server Configuration 16
Locking and Unlocking a Server Configuration 17
Disabling and Enabling a Server Configuration 18
Running a Server Configuration as a Windows Service 18
Deleting a Server Configuration. 20
Setting Server Configuration Options 21
Changing the Logon Sequence Timeout 26
Changing the Inactivity Timeout 27
Changing the Reconnect Timeout 28
Enabling E-mail Support. 28
Enabling E-mail Notification 29
Enabling and Purging the Audit Log 30
Using Pooled Connections. 31
Verifying the Database at Startup (Native-I Vault Only) 32
Changing Maximum Disk Cache Size and Cache Cleanup Interval (Native-I Vault Only) . . 33
Setting the Vault File Lock Time Limit 33
Changing Endpoints 34
Setting Minimum Encryption for TCP/IP 34
Managing Event Handlers 35
Enabling Directory Service Support 37
Diagnosing Server Problems 38
Changing Server Session Options 39
Reviewing Database Information 41

Chapter 4
Working with the Application Logs **43**
Reviewing the Server Log. 43
Copying and Printing Data from the Server Log . . 44
Notifying the System Managers Group 45
Reviewing the Security Event Log 45
Purging Entries from the Security Event Log . . . 46
Reviewing the StarTeam.Log File 47
Customizing the StarTeam.Log File 48

Chapter 5
Setting Up Users and Groups **51**
Understanding the Default Groups 51
Adding Groups. 52
Changing the Parent of a Group. 53
Determining Who Is a Member of a Group 53
Removing an Empty Group 54
Adding Users 55
Licensing. 57
Licensing Users Using Borland License Server . 58
Licensing Users Using StarTeam. 59
Assigning Licenses to Specific Users 60
Changing a User's Password 61
License Server Support 62
Changing a User's Group Membership 63
Removing a User from a Group 64
Checking a User's Logon and Account Status. . . . 65
Removing a User's Account 66
Suspending a User's Account 67
Reactivating a User's Account. 67
Forcing a User to Log Off 68
Forcing a Password Change. 69
Setting the Number of Logon Attempts 71
Reactivating Administrative Accounts 72

Chapter 6
Managing Projects **73**
Before Creating a Project 73
After Creating a Project 74
Adding Access to a Server Configuration 74
Creating a Project 75
Creating a Project by Dragging a Folder. 76
Opening an Existing Project 76
Modifying Project Names or Descriptions 77
Using Keyword Expansion. 77
Forcing Users to Provide a Revision Comment . . . 78
Forcing Users to Lock Files Before Check-In 78
Making Unlocked Files Read-Only. 79
Understanding Process Rules. 79
Requiring the Use of Process Items 80
Using APEs with Projects 81
Deleting Projects. 82
Assigning Project Access Rights 82

Chapter 7
Using Promotion States **83**
How an Organization Might Use Promotion States . 83
Creating Promotion States. 85
Promoting a View Label 85
Demoting a View Label to a Previous State 86

Basing a Configuration on a State 86
Setting Access Rights for Promotions 87

Chapter 8
Using View Manager 89
Preparing to Use View Manager 90
Minimize What Must be Replicated 90
Understand View Manager Capabilities 90
Using the View Manager Main Window 92
Using the Upper Panes 93
Using the Lower Pane 96
Using the View Manager Graphical Interface 100
Overview 100
Selecting Source and Target Objects 101
Configuring the Source View 102
Managing Server Configurations 102
Performing Preview Operations 103
Performing Synchronization Operations 104
Performing Copy-to-Target Operations 105
Observing Progress 106
Resuming an Aborted Operation 107
Using the Log Files 107
Understanding Restart Files 108
Setting and Storing Options 108
Using Properties Files 118
Understanding the Configuration File 120
Finding Objects That Have Changed 120
Comparing and Resolving Text Files 121
Resolving Binary Files 121
Identifying Files Checked in by View Manager . 121
Using the View Manager Command Line 123
Starting View Manager from the Command Line 123
View Manager Command Line Syntax 124
Logging Screen Output 135
Learning from Examples 136

Chapter 9
Using Views 137
Understanding Different Types of Views 137
Branching Views (Branch All, Branch None) . . 138
Reference Views 140
Read-Only Reference Views 140
Non-Derived Views (Blank Branching Views) . . 140
Creating Additional Views 144
Controlling the Branching Behavior of Shared Items 146
Changing a View's Working Folder 146
Deleting Views 147
Assigning View Access Rights 147

Chapter 10
Comparing and Merging Views 149
Understanding View Comparison and viewmerge Differences 150
Using the View Comparison Utility 150
Recording and Tracking Merge Information . . . 151
Displaying a View Comparison Window 152
Changing a View or Its Configuration 153
Using the Folder Hierarchies 153
Comparing Folder Contents Manually 154
Comparing Child Folders 154
Comparing Files 154
Comparing Items Other than Files 155
Controlling the Items That Display 155
Selecting Items from the Item Pane 157
Finding Differences 158
Sharing Items 160
Deleting Items 161
Comparing Item Properties 161
Comparing Files 162
Merging Item Properties Manually 162
Effects of Merging Change Requests on Workflow 163
Merging Files Manually 164
Merging Text Files 165
Merging Binary Files 165
Merging Item Properties Automatically 166
Reconciling Items 166
Reviewing Item References 168
Reviewing Item Properties 168
Refreshing the Panes 168
Merging Views from the Command Line 169
Logging Information 176
viewmerge Examples 177
Troubleshooting 178

Chapter 11
Customizing Item Properties 179
Creating a User-Defined Property Field 180
Creating an Enumerated Type 181
Creating a Numeric Type 181
Creating a Text Type 182
Creating a Date/Time Type 182
Creating a User ID Type 183
Modifying an Existing Field 183
Translating Display Names and Enumerated Values 186

Chapter 12
Controlling Access to Objects 189
Setting Server-level Access Rights 190
Reviewing the Security Log 191
Understanding Access Rights for Projects, Views, Folders, and Items 192
Determining Access Rights and Tokens 192
Using or Ignoring Ownership and Privileges . . . 193
Setting Group Privileges 193
Understanding Access Right Levels 195
Opening Projects and Views 195
Granting Access Rights at the Project Level . . . 196
Granting Access Rights at the View Level 197
Granting Access Rights at the Folder Level . . . 198
Granting Access Rights at the Item Level 199
Automatically Granting Access Rights 200
Denying Access Rights 200
Setting Access Rights 201

Understanding Access Rights for Promotion States . . . 212
Setting Promotion State Access Rights . . . 212
Understanding Access Rights for Components, Filters, and Queries . . . 213
Setting Component-level Access Rights . . . 213
Setting Individual Filter and Query Access Rights . . . 215
Using the Security Guidelines . . . 216
From the Server . . . 216
From the application . . . 217

Chapter 13
Using StarTeam with Visual SourceSafe 219
Getting Ready . . . 219
Understanding Branch Support . . . 220
Wrapping the Application Around VSS . . . 221
Security Issues . . . 223
VSS to StarTeam Interoperability . . . 223
StarTeam to VSS Interoperability . . . 225
Refreshing Data . . . 226
Using StarTeam Exclusively . . . 226

Chapter 14
Using StarTeam with PVCS 229
Understanding the Terminology Differences . . . 229
Understanding PVCS Configuration Files . . . 230
Getting Ready . . . 232
Understanding Branch Support . . . 233
Wrapping StarTeam Around PVCS . . . 234
Security Issues . . . 235
PVCS to StarTeam Interoperability . . . 236
StarTeam to PVCS Interoperability . . . 237
Refreshing Data . . . 238
Operational Refresh . . . 238
Folder Refresh . . . 238
Using StarTeam Exclusively . . . 239

Appendix A
Reporting on Licenses 241
Running the License Report . . . 241
Understanding Licenses . . . 242

Appendix B
Initialization Files 243
Understanding ConnectionManager.ini . . . 244
Understanding starteam-server-configs.xml . . . 244
Understanding starteam-client-options.xml . . . 248

Appendix C
Backing Up the Application 249
What to Back up . . . 249
How to Perform Back ups . . . 250
Online Backups . . . 250
Offline Backups . . . 250

Appendix D
Using the starteamserver command 253
Creating a New Server Configuration . . . 254
Starting a Server Configuration . . . 256
Understanding starteamserver Options . . . 257

Appendix E
Using the Server Administration Utility 265
Upgrading to StarTeam 2005 or Later Releases . . . 266
Creating New Hives . . . 267
Viewing Hive Properties . . . 269
Changing the Archives Path . . . 269
Changing Hive Properties . . . 270
Converting Native-I Vault Archive Files . . . 270
Converting Files as a Background Process . . . 271
Using the ConvertVault Command-Line Application . 271
Converting Files Manually . . . 272
Verifying Tip Revision Information (Native-I Vaults) . . . 272
Purging Deleted Views from a Server Configuration . . . 273
Migrating a Server Configuration to Another Database . . . 274
Exporting Database Information . . . 276

Chapter 15
Using Import/Export Manager 279
Flexible Extraction . . . 279
Server Inaccessibility . . . 279
Installing Import/Export Manager . . . 280
Importing and Exporting . . . 280
Using the Export Utility . . . 281
Copying the Data . . . 281
Using the Import Utility . . . 281
Creating the .cdl Files . . . 282
Mapping Items . . . 284
Setting the Log Properties . . . 286
Understanding Import/Export Manager Limitations . . . 286
Deletions . . . 286
Moved Items . . . 287
Shared Items . . . 287
Illegal XML Characters . . . 287
Known Issues . . . 288
Using Server Scope . . . 288
Using Project Scope . . . 289
Additional Project Scope Export/Import Restrictions . . . 291
Using View Scope . . . 291
Additional View Scope Export/Import Restrictions . . . 293
Using Folder Scope . . . 293
Additional Folder Scope Export/Import Restrictions . . . 296

Appendix F
Entering Change Requests 297
Sample Procedure: How to Enter a Change Request . . . 297
Using Change Request Fields . . . 298
Description Syntax/Shorthand . . . 299

Appendix G
Glossary **301**

Index **311**

Chapter

1

Preface

This manual contains information for StarTeam administrators and project leaders who will be configuring and managing the Server. This administration may be either remotely using a StarTeam client or locally using the Server Administration Tool.

This manual assumes that the administrators are familiar with:

- Creating and modifying relational databases.
- Working with the features of their operating system, such as creating files, running executables, and managing access rights.
- Basic configuration management concepts.

This manual also assumes that the administrators will:

- *Never* modify database contents other than through a client or Server tool.
- *Never* modify vault files other than through a client or a Server tool.

Application Documentation Set

The documentation set is your guide to using the StarTeam product suite. It consists of the following manuals. Depending upon which products you purchased and installed, not all of the applications described in the manuals will be on your system.

- *StarTeam New Features Guide*

  This *Guide* describes the new features in this release.

- *StarTeam Installation Guide*

  This G*uide* contains detailed instructions for installing the core products and lists the system requirements for each of those products.

- StarTeam Getting Started Guide

This G*uide* presents two tutorials on configuration management for the Server and clients. The first tutorial is for administrators and explains how to perform basic administrative functions in both the server and the client. The second tutorial is designed for users and explains how to use basic features in the client.

- *StarTeam User's Guide*

  This *Guide* explains how to use the Windows and Cross-Platform clients to track and manage changes to files, share files and comments with other team members, access previous versions of a file, and other configuration management functions. It also explains how to use items that are not files, such as change requests, requirements, tasks, and topics.

- *StarTeam Administrator's Guide*

  This *Guide* explains how to configure and maintain the Server, create and maintain objects—such as projects and views, and manage user access. This manual also contains instructions on using the Server with existing PVCS and Visual SourceSafe file archives.

- StarTeam Extensions User's Guide

  This *Guide* explains how to design and manage StarTeam Extensions, such as alternate property editors (APEs). It also covers the Workflow Designer and Notification Agent.

- *StarTeamMPX Administrator's Guide*

  This G*uide* explains the basic operation and architecture of a StarTeamMPX system, and presents instructions on installing and configuring its components.

- *StarDisk User's Guide*

  This *Guide* explains how to install and use the StarDisk client application. StarDisk enhances Windows Explorer so you can use it as a simple client for the Server. After you map a specific configuration of a project view to a Windows virtual disk, you can perform basic version control operations on the files and folders within that view.

Online Documentation

The installation procedure installs online versions of the manuals in the \Online Documentation folder of the StarTeam server or client application. The following table lists file names for the online manuals. The online manuals you receive will vary according to the products you purchase.

File Name	Manual
admin.pdf	StarTeam Administrator's Guide
adminMPX.pdf	StarTeamMPX Administrator's Guide
extensions.pdf	StarTeam Extensions User's Guide
install.pdf	StarTeam Installation Guide
newfeatures.pdf	StarTeam New Features Guide
stardisk.pdf	StarDisk User's Guide
start.pdf	StarTeam Getting Started Guide
user.pdf	StarTeam User's Guide

Important The online manuals are distributed in Adobe Acrobat (.PDF) format and require the Adobe Acrobat reader, 4.0 or higher, in order to display them. The installation program for the Adobe Acrobat Reader is located in the \Docs folder of the Installation CD. The reader is also available from the Adobe web site at: www.adobe.com.

Contacting Borland Support

Borland Software Corporation is committed to providing world-class services in the area of consulting and technical support. We have over 15 years of experience in supporting developers and enterprise customers. Our qualified technical support engineers are prepared to handle your support needs on a case-by-case basis or in an ongoing partnership. Borland provides support worldwide, delivering timely, reliable service to ensure every customer's business success.

For more information about Borland's support services, please see our web site at http://support.borland.com.

From the Web site, you can also access many newsgroups on which users exchange information, tips, and techniques. See http://info.borland.com/newsgroups/ for the latest list of free product newsgroups. Also available on the Internet is the Borland Developer Network site at http://community.borland.com. This Borland Community provides access to product specific information, articles, code examples, and news.

When contacting support, be prepared to provide complete information about your environment, the version of the product you are using, and a detailed description of the problem.

For support on third-party tools or documentation, contact the vendor of the tool.

Documentation Conventions

The documentation uses the following conventions.

Select File > Exit	Indicates a menu selection followed by a submenu selection. The greater-than symbol (**>**) separates the commands to be selected from subsequent menus. For example, "Select the File > Exit command" means to select File from the menu bar and then select Exit from the drop-down menu.
`Fixed-Space Font`	Text appearing in Courier font represents information that you need to enter and displays messages from the system.
italics	Syntax appearing in italics represents information that you replace with the names of your files, child folders, etc. Italics are also used for the names of dialogs and books and for emphasis.
Bold	Syntax appearing in bold represents information that you must use exactly as shown.
[ ]	Square brackets surround optional syntax.
|	A vertical bar separates mutually exclusive choices in syntax.
	When this icon appears in the margin, it indicates that the section or procedure applies to the Windows client, but not to at least one of the other clients. If a section or procedure applies to all clients, no icon appears in the margin.
	When this icon appears in the margin, it indicates that the section or procedure applies to the Cross-Platform client, but not to at least one of the other clients. If a section or procedure applies to all clients, no icon appears in the margin

Note	Identifies supplemental information.
Tip	Identifies information on alternative procedures or other helpful but nonessential information.
Important	Identifies information that is essential to the completion of a task.
Caution	Identifies actions that may result in loss of data or procedures that must be followed to ensure that data is *not* lost.

Chapter

2

Administering the Application

This book explains how to use the features for operations normally performed by administrators.

The Server can be accessed from a variety of clients, including the Windows client, Cross-Platform client, StarDisk, or Web Edition. Each client must have a user name and the correct access rights to access the selected server configuration (instance of the Server).

Users with web browsers can access a Web server that runs Microsoft Internet Information Server (IIIS) and on which Web Edition is installed. Then IIS accesses one or more server configurations.

Your company or team may store its data on several server configurations on one or more computers. Any of these configurations can be accessed from a number of clients.

In the next figure, one user running a client on a workstation accesses six projects on four different computers, each of which runs one or more server configurations.

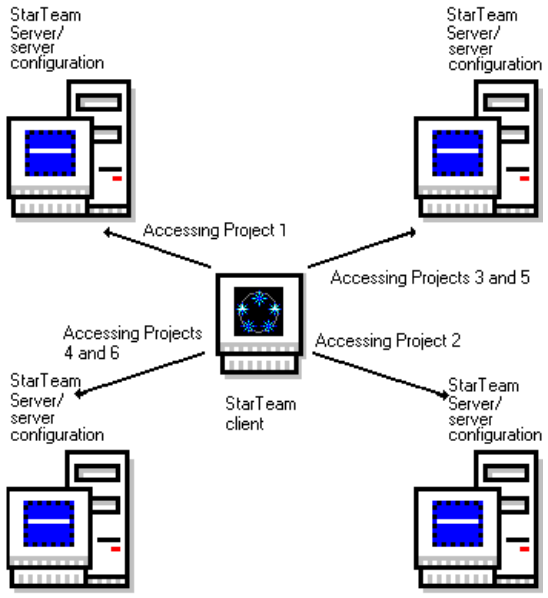

More than one instance of the Server may be running on the same computer. For example, users might run one server configuration with the StarDraw sample project and another with a software development project—both on the same computer.

In the next figure, one user running a client on a workstation accesses projects from two server configurations running on the same computer. Each server configuration has a different name and a different port or endpoint for each protocol. When a configuration is in use, another session using that configuration cannot be started.

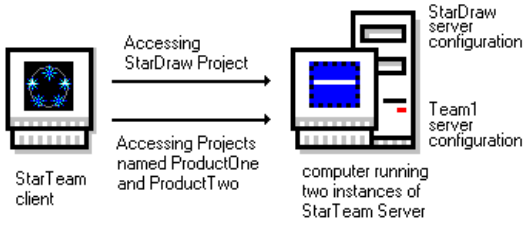

Understanding Server Configurations

Before using the Server, you must decide what database to use and where to store the database and file revisions. Then you must create at least one server configuration.

A server configuration defines a specific database as the repository for its data. To prevent corruption, that database can be associated with only one server configuration. However, that database can be used by other applications.

Any number of projects can be stored in the database associated with a particular server configuration. However, the database must be configured properly to store the amount of data produced by those projects. For more information about specific databases, see the *StarTeam Installation Guide.*

StarTeam 2005 introduced a new vault architecture (Native-II) that provides greater scalability for all server configurations created with StarTeam 2005 or later and for server configurations converted to Native-II vault format with StarTeam 2005 or later. Server configurations now have one or more hives. A hive is a logical disk container of files that includes an Archive area and a Cache area. The Archive area consists of a folder tree in which unique file revisions are stored. The Cache area consists of a folder tree that stores uncompressed file revisions on a temporary basis. Hives can hold an unlimited number of files, providing increased storage capacity.

Other server configuration settings control where, when, how, and by whom the data is accessed. Some initial settings that you provide for the server configuration are properties that are necessary to start the Server. For example, if the user name and password that allow the Server to access the database are not accurate, the Server cannot run. Before starting the server, you can change these properties to meet your requirements.

Understanding This Guide

This guide is written for people who administer StarTeam, which also involves the administration of the application server. It assumes that the basic administrative responsibilities include:

- Creating server configurations using the computer that runs the Server.
- Maintaining server configurations and repositories by:
  - Performing or overseeing regular backups of all data.
  - Checking compliance with the licensing agreement.

- Adding and importing groups and users.
- Creating projects and views of those projects in one or more server configurations.

  Views provide a variety of ways to look at projects, and ways to branch projects from one product release to another.

  Projects that interoperate with Visual Source Safe (VSS) files cannot use views.
- Maintaining projects and views, including:
  - Setting access rights to control what users and groups can access what objects.
  - Setting up promotion states and controlling the transition from one state to another.
  - Merging some views into one another or back into the root (baseline) view.
  - Creating custom fields (if your company has purchased a license that allows repository customization).
- Setting guidelines for how to use the application.

Chapter

3

Configuring and Running the Server

A server configuration identifies a repository for projects. Each server configuration acquires its own set of projects as they are created.

The Server can run any number of server configurations. However, if more than one configuration runs at the same time on the same computer, each server configuration must use a different endpoint (port). Server configurations can be started manually or run as Windows services.

You must create at least one server configuration before you can run the Server. Because each server configuration must use a database, the process of creating a server configuration depends upon the database selected. Detailed instructions for creating server configurations are presented in the *StarTeam Installation Guide*, which includes a chapter for each type of database.

Planning Server Configurations

Before you create projects and views in a server configuration, you need to be aware that:

- A server configuration can store multiple projects. Each project has one root view, and each view can have any hierarchy of views below it.

The following diagram illustrates these principles:

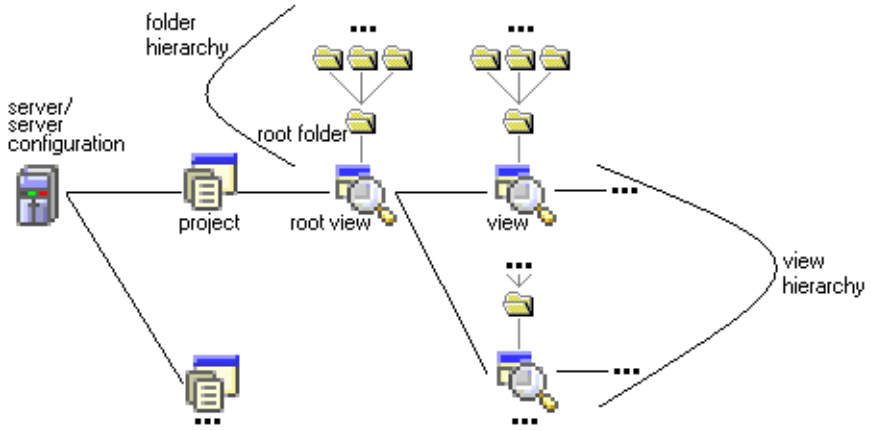

Each view has a single folder as its root folder. The root folder can have any number of child folders. Together, the root folder and its children comprise the folder hierarchy.

Each folder has a designated working folder, which is the default location in which files are stored when they are checked out to your workstation. Working folders are often hierarchical, but they do not have to be. The physical path to a working folder may have little or no correlation with the logical path of the folder.

- You can share projects, views, items, and even folders, as long as their original location and their shared locations are controlled by the same server configuration.
- Promotion states apply only to the view for which they are created.
- View labels (such as build labels) are automatically attached to every item in the view for which the label was created. As a result, you must think about what items will require build labels, put items that require the same build label in the same view, and put items that require different build labels in different views.

Projects and views can be organized in various ways. For example, some teams put each component of a software product in its own project. The projects are placed in the same server configuration so that the components s can share files. Then the teams create views within each project for different releases of each component.

Other teams put all components of a product in the same project. They use views to separate components from one another and to separate different releases of each component.

Some teams have no need to separate components into projects or views, because all the components are built together and receive the same build label. These teams put all components in the same project and use views to separate each release of the complete product.

For more information about projects and views, see Chapter 6, “Managing Projects,” on page 73 and Chapter 9, “Using Views,” on page 137.

Managing Server Configurations with the Server Administration Utility

The Server Administration utility is a Java application that enables administrators to create and manage server configurations and the repositories they access. It is automatically installed with the Server and can be run only from a computer on which the Server resides. From the server, this tool can administer both local and remote configurations, as it can access the server's starteam-server-configs.xml file.

This utility can also be installed with the client, if you choose a custom installation. However, from the client, the tool can administer remote server configuration only.

With the Server Administration utility, an administrator can perform all operations on either remote or local server configurations, including the following:

- Create, enable, disable, or delete a server configuration.
- Display or modify the session options for a server configuration.
- Start or shut down a server configuration.
- Set or remove a server configuration as a Windows service.
- Review the status and execution mode of all server configurations running on this computer.
- Access the Hive Manager.
- Convert or stop the conversion of a Native-I vault to a Native-II vault.

The following tasks can also be performed on remote server configurations from clients on which the Server Administration utility is installed:

- Log onto a server as a different user.
- Add and manage user accounts.
- Set the security policy for a server configuration.
- Assign access rights to users and groups for a server configuration.
- Add, modify, or delete connections to a server configuration.
- Set or modify the configuration options for a server configuration.
- Display the server log file (Server.*locale*.Log).
- Lock or unlock a server configuration.

See Appendix E, "Using the Server Administration Utility," on page 265 for more information.

Starting the Server Administration Utility

Before you use the Server Administration utility to work on a server configuration, you must have administrative privileges for that configuration and the configuration must be running. You can start the Server Administration utility from the command prompt or from the Start menu on Windows systems.

The Server Administration utility can be used to manage server configurations running on the computer on which it is installed or multiple computers running the Server. Connection information for server configurations is stored in the starteam-servers.xml file.

To start the Server Administration utility from the command prompt:

1. Open a command prompt window.
2. Change directories to the Server folder. For example,
   ```
   cd C:\Program Files\Borland\StarTeam Server x.x
   ```
3. Enter the following at the command line:
   ```
   serveradmin
   ```
   The *Server Administration* utility appears.

To start the Server Administration utility from the Windows Start menu:

- From the server, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions run the AdminTool.stjava file. The *Server Administration* dialog appears. The *Server Administration* dialog on the client is similar to that of the Server Administration utility but can be used to administer remote servers only. Some functions, such as migrating a database, can be performed only when a server is shut down.

Understanding the starteam-servers.xml File

The starteam-servers.xml file contains connection information for the server configurations that have been defined for the computer on which the file resides. The location of the file on your computer depends on the type of operating system.

For additional information, see Appendix B, “Initialization Files,” on page 243.

Creating a Server Configuration

When you use the Server Administration utility with StarTeam 2005 or a later release to create a new server configuration, this creates a Native-II vault for file storage. The architecture of this vault format differs from that of the Native-I vaults used with StarTeam 6.0 and earlier. Native-II vaults can handle larger files, provide better performance, and can spread file storage over multiple volumes.

Server configurations created with StarTeam 2005 or a later release require at least one hive (default or custom) and automatically set up a Native-II vault. All files are added to, checked out, and checked into this vault.

Note: Server configurations created with StarTeam 6.0 or earlier releases must be upgraded so that they will have Native-II vaults, into which new files must be placed, as well as their pre-existing Native-I vaults. To upgrade, display the Server Administration utility, select the server, and click Upgrade. For details, see the *StarTeam Installation Guide*.

Archive files in an upgraded server configuration will remain in the Native-I vault until you check them out and check them back in again or until they are converted to Native-II vault format. For details, see “Creating New Hives” on page 267.

Before creating a new server configuration, you need to decide upon a unique name for the configuration. This name is case insensitive and cannot contain colons (:), back slashes (\), or forward slashes (/), but can contain blanks or apostrophes (').

You must also set up the database to be used with the server configuration. A database can contain only one server configuration; however, other applications can share a database with StarTeam.

StarTeam Server places server log files in the location designated as the server configuration's repository path. When you first start a new server configuration, the Server creates the Attachments folder, Hive Index, and other files and folders in the same location. These objects are maintained by the Server; do not delete them. Paths to some of the folders can be changed with the *Server Administration* dialog.

For new server configurations, the initial hive is created at the same time as the server configuration. The user supplies an archive path and a cache path. The default paths are repository_path\DefaultHive\Archives and repository_path\DefaultHive\Cache. If desired, the location of these paths can be changed by using the *Hive Manager* dialog. All new files are added to a hive, even if the existing files have not been converted and remain in the Native-I format.

A single server configuration can have several hives, each of which has its own archive and cache path.

To create a server configuration:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. Click the New Configuration icon or select Server > New Configuration from the menu. This action displays the *Create a New Configuration* dialog.

3. Enter a unique name in the Configuration name text box. For this exercise, enter `StarDraw`.

4. In the text box, enter or browse for the location in which the Server will create the server configuration files.

5. Select a database type from the list box. The selections include Microsoft SQL Server/MSDE, DB2, and Oracle. The database type cannot be changed once the server configuration has been created.

6. Select or deselect the Create new StarTeam database and ODBC data source checkbox. The default is selected. Note: This option cannot be selected for DB2.

7. In the Initial Hive Settings, select the Default or Custom hive option. Note: Only a Custom Hive can be selected for DB2.

   - If you select the Default hive, changing the repository path changes the default hive settings. Changing the repository path does not have this effect if you select a Custom hive.
   - If you are creating a Custom hive, you can override the default hive settings. You can change any of the following fields:
     - Name: Unique name for the hive. DefaultHive is the default.
     - Archive path: Path to the hive's Archives directory. The default is <repository_path\DefaultHive\Archives.
     - Cache path: Path to the hive's Cache directory. The default is <repository_path\DefaultHive\Cache.
     - Maximum cache size: Maximum number of megabytes of hard disk space that the Cache can use. The default is 20% of the disk space available when the option is set.
     - Cache cleanup interval: Seconds between cache cleanup/refresh operations. The default value is 600. The range is 60 (1 minute) to 3153600 (1 year).
     - Storage limit threshold: Percentage of total disk space allowed for hive. When this percentage has been reached, no more archives can be added to the hive. -The default represents 95% of total disk space.

8. When the information is complete, click Next. The information that must be entered on the next screen varies according to the database selected. The ODBC data source cannot be changed after the server configuration has been created.
   - A Microsoft SQL Server or MSDE database requires:
     - Host name
     - Sys Admin (sa) password
     - ODBC data source name
     - New database name
     - New database login name
     - New database password
     - Confirm database password
   - For an Oracle database, enter text in:
     - TNS service name
     - System password
     - New ODBC datasource name
     - New schema user name
     - New schema password
     - Confirm schema password
   - For a DB2 database, enter text in:
     - ODBC data source name
     - Database login name
     - Database password
9. When the information is complete, click Next (for a Microsoft SQL Server or MSDE database) or Finish (for a DB2 database).
   a. For a Microsoft SQL Server or MSDE database, the *Create a Microsoft SQL Server/MSDE Data Source* dialog appears.
      1. Review the dialog.
      2. If you want to edit the size or location of the data and transaction files, deselect Use default configuration and make the changes. Microsoft limits the size of a MSDE database, by license, to 2048 MB. If you require a larger database, you must purchase a license for Microsoft SQL Server.
      3. When the information is correct, check Finish. A message displays to indicate that the server configuration has been created successfully.
   b. For an Oracle database, the *Create an Oracle Data Source* dialog appears.
      1. Review the dialog.
      2. If desired, edit the Tablespace name, Initial Size, and Location.
      3. When the information is correct, close the window. A message displays to indicate that the server configuration has been created successfully.
10. After the new server configuration has been created, it appears in the *Server Administration* dialog with the status of New. Select the configuration, and click Start Server. The Server then initializes the database and creates the server configuration's files and folders. For more details on this process, see "Starting a Server Configuration" on page 15.

The initialization process may take a few minutes. When the Server finishes this activity, the Status column for the server configuration changes from New to Running.

11 After the server configuration starts running, click Exit to close the dialog.

When you create a server configuration, you set values for the configuration session options, which are stored in the starteam-server-configs.xml file. These options specify the core information that the Server requires to start a server configuration.

After the server configuration has been created, you can modify the default server configuration options, which enable you to fine-tune server configuration performance. See "Setting Server Configuration Options" on page 21 for detailed information on these options.

Starting a Server Configuration

You can start a server configuration from the Server Administration utility or from the command prompt with the starteamserver command. This section presents instructions on using the Server Administration utility to start a server configuration. For information about using the starteamserver command, see Appendix D, "Using the starteamserver command," on page 253.

You can also run the server configuration as a Windows service. See "Running a Server Configuration as a Windows Service" on page 18 for more information.

To start a server configuration:

1 From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

This action displays the *Server Administration* dialog.

2 Select the server configuration that you want to start. If you have not yet logged in, you will be asked to do so.

Note: If a server configuration cannot be run with this release of the StarTeam Server, a special icon appears immediately beside the server configuration name. The icon indicates that the server configuration database has been created with Access. Access databases are no longer supported for StarTeam Server upgrades.

3 Do one of the following.

- To start the server configuration with its current configuration options, select the Start Server icon or click Actions > Start Server on the menu.
- To start the server configuration with different configuration options, select the Start with Override... icon or click Actions > Start with Override.

4 Depending upon the server configuration selected, one of the following events will occur:

- If the server configuration was created in StarTeam 2005 or a later release or if its files have been fully converted to a Native-II format, the system continues the startup operation.
- If the server configuration has one or more hives defined, but contains files that are still in Native-I format and conversion is turned off, a message box states: "There are files in this server configuration that have not been converted to the Native-II format and conversion is turned off. Would you like conversion turned on?" If you see this message, click Yes or No. The system then continues its startup operation.
- If the server configuration has no associated hives (that is, it is totally in Native-I format), a message box states: "This server configuration has not been converted to the Native-II format. If you start the server now, a default hive will be

created because all new files must be placed in a Native-II vault hive. Click OK to start the server (and use a default hive) or Click Cancel to create your own hive using the Hive Manager." If you click OK, the system creates the hive and then continues its startup operation. If you select Cancel, you will exit to the *Server Administration* dialog, from which you can access the *Hive Manage*r dialog and create a new hive. You can then return to the *Server Administration* dialog and start the server configuration.

5 If you initially selected the Start with Override option, the *Start with Override* dialog appears.

6 Modify the fields as appropriate, and then click OK. The server configuration information will then be updated.

If you are already using the default endpoint (49201) for another server configuration, the first time you start a new configuration you may want to use an override for the endpoint. This action sets the endpoint to the one that you will want to use in the future.

7 After the server configuration finishes its startup procedure, its Status column changes from Ready to Running.

The first time you start a new server configuration, the Server performs several startup tasks. It creates and initializes the database to be used by the server configuration, installs the stored procedures for that database type, and creates the repository folders and the hive used by the configuration. This process may take several minutes.

Logging on as a Different User

Sometimes a user has more than one user name. For example, a QA team leader may need to log on as an individual and as the QA team leader.

If you are already logged onto a server configuration or are running the application Toolbar, but wish to log on as a different user, you can do so.

- If you are already logged on, the user name you used most recently appears in parentheses after the server name in the *New Project Wizard* and *Open Project Wizard* dialogs.
- On the Toolbar, the user name in shown parentheses is the one recognized as your default set of credentials.

Note The server must be running to perform this operation.

To log on as a different user:

1 On the *Server Administration* dialog, select the server configuration to be accessed.

2 Click Actions > Logon As. The *Log On To* dialog box opens.

3 Enter the alternate user name and password in the appropriate text boxes.

4 (Optional) To reset your default credentials for this server configuration to the user name and password you just entered, select the “Save as default credentials for this server” check box.

5 Click OK.

Shutting Down a Server Configuration

Although you do not need to shut down a server configuration created with StarTeam 2005 or a later release to perform a backup (unless it contains files that have not yet been converted to Native-II format), you may need to do so to perform other maintenance tasks.You can perform a shutdown from:

- The *Server Administration* dialog.
- The command line. For information about using the command line, see Appendix D, "Using the starteamserver command," on page 253.

Note for Enterprise Advantage Users

If you are running the Server as a service and the Notification Agent as a dependent service, you cannot shut the Server down unless the Notification Agent service is shut down first.

To shut down a server configuration:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (if installed).

   This action displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to shut down. If you are performing this task from a client, you must select a remote server.
3. From the *Server Administration* dialog, click the Shut Down Server icon or select Actions > Shut Down Server. The system displays the *Server Shutdown* dialog.
4. The system displays a message asking you to confirm that you want to shut down the server configuration. Click OK to confirm this decision.

Locking and Unlocking a Server Configuration

Locking a server configuration enables you to limit access to that configuration while you perform backup procedures or database maintenance. When a server configuration is locked, only server administration commands are accepted. For any other command—such as checking out files, the Server sends an exception message stating that the server configuration is unavailable.

You can also lock and unlock the server from the Server Administration utility or from the command line, using the stcmd server-mode command. See the *StarTeam User's Guide* for more information about the command line.

Note The server must be running to perform these operations.

To lock or unlock a server configuration:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. From the list of servers, select the server configuration you want to lock or unlock. If you are using the client, you can administer remote configurations only.
3. Do one of the following:
   - Click the Lock Server icon or select Actions > Lock Server. On the resulting dialog, indicate whether you want to:
     - Lock the Server. This option allows minimal administrative options, primarily, start, shutdown, lock and unlock operations. It is usually done for backup operations in environments where server activity is not 24/7.

- Lock the Server for Exclusive Use by <user name>. This option, which displays the user's logon name, allows a user to lock the server for his or her use only.

- Click the Unlock Server icon or select Actions >Unlock Server, to enable the server-level operations.

Disabling and Enabling a Server Configuration

You can disable or enable a server configuration from the Server Administration utility. Disabling a server configuration enables you to take a server configuration "out of service" and ensure that it is not started by accident. For example, if you migrate a server configuration, you should disable the prior server configuration. After you are sure that the new server configuration and database are working properly, you can delete the prior server configuration. You can also re-activate a disabled server configuration.

Note Only a server that is shut down can be disabled or enabled.

To disable or enable a server configuration:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. On the *Server Administration* dialog, select the server configuration you want to disable or enable.

3. Click the Disable Local Server or Enable Local Server icon or select Server > Actions > Disable/Enable Local Server.

   - If the server configuration is currently enabled (i.e., has a "Ready" status), it becomes disabled.

   - If the server configuration is currently disabled, it becomes re-enabled.

Running a Server Configuration as a Windows Service

Under Windows, a server configuration can run as a service, which gives it operating system priority.

By default, a Windows service:

- Starts automatically when you start the computer on which the Server runs.

- Logs on as a system account.

- Uses the server configuration specified at the time the service was set.

If a server configuration is newly created, you must run it once before it can be set as a service.

Even when a server configuration is running as a service, you can shut it down from the *Server Administration* dialog or the command line. The Windows service will still run, but the Event Viewer Application log and the Server's server log will show that the server configuration has been shut down.

If a server configuration is set as a Windows service but the service has been stopped or the server configuration has been shut down, you can restart the server configuration by using the Server Administration utility or the command line. This action does not restart the Windows service, but simply restarts the server configuration.

To set a server configuration to run as a service:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. Select the server configuration you want to set as a service.
3. If the server configuration is running, click the Shut Down Server icon or select Actions > Shut Down Server.
4. Click the Set/Remove Service icon on the toolbar or select Server > Set Service from the menu. The *Log On Service As* dialog appears.
5. Do one of the following to select the user name for the service:
   a. To use the local system account, select the Local System Account check box.
   b. To use a specific user's account:
      1. Clear the Local System Account check box.
      2. Enter an account name. The usual format is *DomainName**UserName*. If the account belongs to a built-in domain, you can use .*UserName*.
6. Click OK.

These actions update the Execution Mode column for that server configuration to "Service." The next time you start the server configuration or restart your computer, the server configuration will run as a service.

To determine whether a server configuration is running as a Windows service, locate the server name in the left panel of the *Server Administration* dialog. Beside the name, an icon indicates whether the server is enabled and/or running as a service.

If you want to discontinue running a server configuration as a service, you must first stop the server configuration, and then remove the service using the Server Administration.

To stop running a server configuration as a service:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. Select the server configuration you want to remove as a service. If you have not logged in yet, you will be asked to do so.
3. If the server configuration is running, click the Shut Down Server icon or select Actions > Shut Down Service.
4. Click the Set/Remove Service icon or select Server > Remove Service.

The server configuration will no longer run as a service. The Execution Mode column for the server configuration will be updated to "Application."

If a server configuration that is set as a service fails for any reason or has been shut down, Windows records that information in the Event Viewer Application log.

To display the Event Viewer Application log:

1. From the computer on which the Server is installed, select Start > Programs > Administrative Tools (Common) > Event Viewer. The *Event Viewer* dialog appears.
2. Select Log > Application from the Event Viewer menu bar. The *Event Viewer - Application Log* dialog appears, and the server configuration information displays.
3. Double-click the log entry to view the *Event Detail* dialog.

Deleting a Server Configuration

When you delete a local server configuration, the information about that server configuration is removed from the starteam-server-configs.xml. After this information is gone, the Server can no longer start or access the server configuration. However, the contents and location of the files and database associated with the server configuration are unchanged. You must manually delete them.

You can delete a server configuration from the Server Administration utility or the command prompt. A server cannot be deleted while it is running.

When you delete a remote server configuration, this merely eliminates access to that server configuration by removing the server configuration from the list of servers. The server configuration can be re-added if appropriate. It is *not* a permanent deletion.

To delete a server configuration:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. On the *Server Administration* dialog, select the server configuration you want to delete.

3. If the server configuration you want to delete is running, click the Shut Down Server icon or Actions > Shut Down Server.

4. The system displays a message asking you to confirm that you want to shut down the selected server configuration.

5. Click OK to confirm that you wish to shut down the server.

6. Click the Delete Server icon or select Server > Delete Server.

   The system displays a message asking you to confirm that you want to delete the selected server configuration.

7. Click OK to confirm that you wish to delete the server.

   The system displays a message informing you that the deletion was successful and alerting you that the repository, hive, database, and database connection (the ODBC DSN) for the server configuration have not been deleted. These items can only be deleted manually.

8. Click OK.

To delete a server configuration using the starteamserver command:

1. If necessary, log onto the computer on which the Server is installed.

2. Open a command prompt window and navigate to the server's installation folder.

3. Ensure that the server configuration you want to delete is not currently running. At the command line, enter:

   ```
   starteamserver -list
   ```

   The system displays a list of all server configurations and their status: Ready, Starting, Running, Disabled, or Stopping.

4. If the server configuration you plan to delete is running, stop the server by typing the following at the command line:

   ```
   starteamserver -stop ConfigurationName
   ```

   The system displays the following.

   ```
   Server ConfigurationName shutting down...
   Server ConfigurationName shut down successfully.
   ```

5. Delete the server configuration by typing the following command:

```
starteamserver -remove ConfigurationName
```

The system displays the following.

```
Configuration ConfigurationName removed successfully.
```

Setting Server Configuration Options

The Server stores server configuration options in the database you specified when the configuration was created. The first time a server configuration is started, the Server automatically adds the server configuration options, with default values, to the database. You should review these options, because you may not want the default settings. For example, a server configuration does not enable e-mail by default, but many administrators choose to use this feature.

Server configuration options must be set while the server is running.

To set server configuration options:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. In the list of servers, select the server configuration you want to access. If you are using the client, you can administer remote clients only.

   If the server configuration does not yet appear in the list, you can add it as follows:

   a. Click the Add Server icon or select Server > Add Server from the menu. The *Add Server* dialog appears.

   b. In the Server Description text box, enter a descriptive name for the server configuration.

   c. In the Server Address text box, enter the computer name or IP address.

   d. In the TCP/IP Endpoint text box, enter the correct endpoint. The default endpoint is 49201.

   Note If you specify an invalid endpoint or one for a server that is not running, the Server returns an error.

   Server configurations running on the same computer must use different endpoints (ports). You can create more than one server configuration with the same name and endpoint. However, if you attempt to run these configurations at the same time, only the first one selected will start successfully. The remaining server configurations appear to start but are ignored by the Server.

   You can change a server's endpoint at any time by restarting the server configuration with an override. See "To start a server configuration:" on page 15 for more information on using overrides.

   e. (Optional) If desired, select Compress Transferred Data.

   f. (Optional) Select the level of encryption desired for this server configuration.

   g. Click OK. The server configuration appears in the list of servers in the *Server Administration* dialog.
3. From the *Server Administration* dialog, click the Configure Server shortcut or select Tools > Administration > Configure Server. If you have not yet logged onto the server configuration, you will be asked to do so. After you log on, the *Configure Server* dialog appears.

4 Edit the options on the *Configuration* Server tabs as appropriate.

Table 3.1, “Server Configuration Options” lists each option on a tab-by-tab basis.

5 Click OK.

6 Depending upon the options you have modified, the Server may display the following message:

```
Server configuration changed. The server must be restarted for the changes to
take effect.
```

To restart a Server from a command line, enter the following:

```
starteamserver -restart ConfigurationName
```

To restart a Server from the *Server Administration* dialog (when you are working from the Server Administration utility on the server):

1. Click the Shut Down Server icon or select Actions > Shut down Server.
2. On the resulting *Server Shutdown* dialog, select Shut Down And Restart Selected Server and click OK.

   If you have changed the endpoint, click Shut Down Selected Server on the *Server Shutdown* dialog. Then start the server configuration with either the starteamserver command or the *Server Administration* dialog.
3. Click OK when the Server displays a message asking you to confirm the command.

The following table lists each server configuration option on a tab-by-tab basis.

Table 3.1 Server Configuration Options

Options	Default	Comments
General Tab		
Server startup log file	..\Repository Path\ server.log	Read-only; path specified when creating a new server configuration.
Attachments path	..\Repository Path\ Attachments	Editable path; folder created by the Server.
Logon sequence timeout	60 seconds	Any logon not completed within this amount of time will fail.
Inactivity timeout __ minutes	Off	Automatically logs off users who are inactive for the specified amount of time. Does not apply to users who have set system notification in their Personal Options for a shorter period of time, because of the automatic communication between the client and the server. Also does not apply to named users, if the Except Named Users option (shown below) has been selected.
Exclude named users	Disabled	Allows named users to remain logged on even if they have exceeded the Inactivity Timeout limit. Feature is available only when Inactivity Timeout is selected and a value entered.
Reconnect timeout __ minutes	30 minutes	Determines the amount of time the client has to reestablish a lost connection with the server. Automatic reconnection cannot be performed if the server has been restarted.
Enable e-mail support	Off	Allows users to e-mail items to other users from within the application, even when the recipients are not running the application. This feature must be enabled to select the e-mail notifications option. When e-mail support is enabled, an e-mail address must be entered for each user.
SMTP server	Disabled	Required if e-mail enabled.
Port	Disabled	Default SMTP port is 25 if e-mail is enabled.

Table 3.1 Server Configuration Options (continued)

Options	Default	Comments
Enable e-mail notification	Off	Available when Enable e-mail support is selected, an SMTP server is enabled, and a Port for the SMTP server is specified. If notification is enabled, a team member will receive e-mail when a change request becomes that person's responsibility, when any field for a requirement or a task changes and the team member is responsible for that requirement or task, or if any field for a topic changes and the team member is listed as a recipient for that topic.
Audits Tab		
Enable audit generation	On	Audit log data is stored in the database for the server configuration; if data requires too much space, option can be disabled.
Purge audit entries older than __ days	Off	Automatically removes audit entries older than a specified number of days to minimize the amount of log space required. Default is 90 days, if option is enabled. Number of days can be edited.

Table 3.1 Server Configuration Options (continued)

Options	Default	Comments
Database Tab		
Database type	Disabled	Read only; database type can be set only when server configuration is created.
DSN	Disabled	Read only; item can be set only when server configuration is created.
Pooled connections	10	Increasing the number of pooled connections can improve database throughput. Borland recommends a setting of 20-25 with Microsoft SQL Server.
Native-I Vault Tab (Does not appear for Native-II Vaults)		
Vault path	..\Repository Path\Vault	Editable path; folder created by the Server.
Disk cache path	..\Repository Path\Vault\ Cache	Editable path; folder created by the Server.
Archive path	..\Repository Path\Vault\ Archive	Editable path; folder created by the Server.
Verify database at startup	Off	Borland recommends verifying the database *only* after restoring a backup or when experiencing database problems; for more details, see page 32. Applies only to Native-I vaults.
Attempt automatic recovery	Disabled/Off	Available when Verify database at startup option is selected; if enabled, the verification utility reconciles differences between the server configuration's database and the archive.
Stop server on errors	Disabled/Off	Option that causes the server to stop if the verification utility reports errors; available when Verify database at startup option is selected.
Maximum cache size	100 MB	Borland recommends starting with default size, then editing as required.
Vault file lock time limit	60 seconds	Borland recommends starting with default size, then editing as required.
Cache cleanup interval	10 seconds	Borland recommends starting with default size, then editing as required.
Protocol Tab		
TCP/IP endpoint	49201	Selected during creation of server configuration.
TCP/IP encryption levels	No encryption	Used to set a minimum encryption level for data transferred via TCP/IP; use **Add**, **Remove**, and **Modify** buttons to add additional encryption levels.
Event Handlers Tab		
Event handler	None	Allows entry or selection of event handler program.
Event handler description	On	Allows description of selected event handler program.

Table 3.1 Server Configuration Options (continued)

Options	Default	Comments
Directory Service Tab		
Enable directory service support	Off‘	Uses the specified Microsoft Active Directory service to validate user passwords. For a user's password to be validated, the Validate with directory service option must also be selected on the New User Properties or User Properties dialog and the Distinguished Name from Microsoft Active Directory service entered for the individual. Restart the StarTeam server configuration to be sure that the connection to the service can be made before setting up the users. The server log contains the connection information; for example, "Connected to Active Directory Server: ldaps://host:port" where host and port are the values you enter on this tab.
Host	None	Host name or IP address of the Microsoft Active Directory service; alphanumeric value of up to 254 characters. Instead of using a host name or IP address in the Host text box, you can use a domain name. When you use a domain name, StarTeam Server can contact any active copy of Active Directory anywhere in the domain so long as that copy uses the specified port. Some companies run more than one copy of Active Directory in case one goes down.
Port	636 (secure port)	Secure sockets layer port of the directory server; numeric value.
Use as secure port	On	Indicates whether the port is secure (default) or non-secure.
Diagnostics Tab		
Trace operations that take at least ___ milliseconds	0 milliseconds	Creates a .trc file that allows commands to be traced. Commands are traced if they have a duration time that equals or exceeds the specified number of milliseconds. If 0 (the default) is used, all commands will be traced.
Unexpected conditions	Off	Creates a diagnostic (.dmp) file for asserts (server log entries with code # 8).
Errors	Off	Creates a diagnostic (.dmp) file for exceptions (server log entries with code #4).

Changing the Logon Sequence Timeout

The logon sequence timeout setting applies to both a client and the server configuration. This is the amount of time the client has to make the connection to the Server. If this time expires and a connection was not made, the user must try to log on again.

This operation can be performed only when the server is running.

To change the logon sequence time:

1 Do one of the following:

- From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
- From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 From the list of servers, select the server configuration that you want to change.If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.

3 Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.

4 Select the General tab.

5 In the Logon Sequence Timeout text box, enter the number of seconds users have to log on. The maximum logon sequence time is five minutes.

6 Click OK.

Changing the Inactivity Timeout

The inactivity timeout is a security feature that automatically logs users off when they have been inactive for the length of time specified by the administrator. If a client has no communication (either automatic or manual) with the server configuration for that length of time, the Server drops the connection and the user's session is deleted from the server. If the user has a concurrent license, that license is automatically returned to the pool of concurrent licenses.

To allow named users (that is, users with a fixed license) to remain logged on, even if they exceed the inactivity timeout limit, administrators can select the Exclude Named Users option after selecting the Inactivity Timeout option and entering a timeout value.

Even if an inactivity timeout value is set, users will not time out if their system notifications are set for a period of time that is shorter than the inactivity timeout. For example, suppose a user has notification set to automatically check for new change requests every ten minutes and the inactivity timeout is set for 60 minutes. In this case, because of automatic communication between the client and the server, the user will never time out.

This operation can be performed only when the server is running.

To set an inactivity timeout for users:

1 Do one of the following:

- From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
- From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.

3 Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.

4 Select the General tab.

5 Select the Inactivity Timeout check box.

6 Enter the number of minutes in the Inactivity Timeout text box.

7 (Optional) If you want to allow named users (that is, users with a fixed license) to remain logged on, even when they exceed the Inactivity Timeout limit, select the Exclude Named Users option.

8 Click OK.

Changing the Reconnect Timeout

If a client loses its network connection, users are disconnected from the server. The reconnect timeout option determines the amount of time the client has to re-establish the connection. The client attempts to reconnect only if the user is trying to send a command to the server. A reestablished connection contains the full context of the lost connection.

If the client successfully reestablishes its connection to the server within the window of time set in the Reconnect Timeout, users can simply continue working in the application. They do not have to close their projects, log in again, and reestablish their view settings.

The reconnect timeout can be changed only on a server that is running. It does not work when the server has been restarted.

Note: When a server must be restarted, the client cannot automatically reconnect to the server. Also, if the Reconnect and the Inactivity Timeout are both enabled and the Inactivity Timeout is longer, the user may be logged off before the client can re-establish the connection.

To change the reconnect timeout :

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.
3. Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the General tab.
5. Select the Reconnect Timeout check box to enable this option.
6. Enter the number of minutes in the text box to set the reconnect timeout value. The default time is 30 minutes.
7. Click OK.

Enabling E-mail Support

When e-mail is enabled for a server configuration, users can e-mail an item's properties to another user from within the application. The e-mail recipients do not need to be running the application to receive the e-mail. See the *StarTeam User's Guide* for more information about mailing item properties.

Automatic e-mail is sent to users when their exclusive locks on items are broken. Users can only break locks if they have the correct access rights and privileges to do so. See "Setting File Access Rights" on page 207 and "Setting Access Rights for Other Items" on page 209.

You can also configure the application to perform automatic e-mail notification when certain other events occur. Depending on the server configuration and system policy options you select:

- Members of the System Managers group can receive e-mail whenever an error is added to the server log. See "Notifying the System Managers Group" on page 45.

- Members of the Security Administrators group can receive e-mail whenever a logon failure occurs. See "Setting the Number of Logon Attempts" on page 71.
- All users can receive automatic notifications about items for which they are responsible or for which they are recipients. See "Enabling E-mail Notification" on page 29.

E-mail support can be selected only on a server that is running.

Note If a recipient of an item or notification has an incorrectly formatted e-mail address, an entry is written to the server log indicating that there was a problem sending e-mail to that address. If an e-mail address is formatted correctly but is invalid (as in "junk@place.com"), the e-mail is sent to all valid recipients, and the sender gets an "Undeliverable message" from the e-mail system for the invalid address.

To enable e-mail support:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.
3. Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the General tab.
5. Select the Enable E-mail Support check box.
6. In the STMP Server text box, enter the host name for your SMTP server. You can use an IP address if your site uses only static IP addresses.

   The application uses the SMTP (Simple Mail Transfer Protocol), which traditionally operates over TCP using port 25. It is widely used and is the Internet's standard host-to-host mail transport protocol.

   For Windows environments, the Exchange server is usually the SMTP server.
7. (Optional) The Port text box default value is 25. Change this value if your SMTP server uses another port.
8. Click OK.

Enabling E-mail Notification

If you enable e-mail notification, a user will automatically receive e-mail:

- If the user is assigned the responsibility for a change request.
- If any field for a requirement or task for which the user is responsible has changed.
- If any field for a topic for which a user is listed as a recipient has changed. (If no recipients are listed for a topic, no one receives notification.)

Team members do not have to run a client to receive notification messages, because e-mail notification is client-independent.

For requirements, task, and topics, the Server sends the contents of all the non-advanced fields for the changed item, omitting any fields that are empty. Changed fields have a series of four asterisks after their names.

Messages sent to recipients of automatic e-mail notification are localized, based on the server's locale. For example, if the server's locale is fr-FR, the message is sent in French because the Server has been localized in French. When no translation is available for a locale, the message is in English.

The language used with a specific server configuration can be changed by adding NotificationLocale to the section of the starteam-server-configs.xml reserved for the configuration. For example, if you add NotificationLocale=ja, the messages are sent in Japanese.

When users are running a client, they can also be notified about items for which they are responsible or are recipients by selecting the notification option from the *Personal Options* dialog. When this option is selected, the appropriate icons display on each workstation's status bar (in the system tray).

Users may confuse e-mail messages sent by individuals with e-mail notification messages, because they are somewhat similar. Therefore, it is a good idea to let users know when you enable automatic e-mail notification and to explain the differences between the two types of e-mail messages and the two types of notification. For more information, see the *StarTeam User's Guide*.

E-mail notification can be enabled or disabled only on a server that is running.

To enable/disable e-mail notification:

1. You must enable e-mail from the General tab of the *StarTeam Server Configuration* dialog if you want to use e-mail notification. See "Enabling E-mail Support" on page 28 for details.
2. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
3. From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.
4. Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.
5. Select the Notifications tab.
6. Do one of the following:
   - Select the Enable E-mail Notification check box, so that all users will receive e-mail notifications about changes.
   - Clear the check box, so that no users receive e-mail notification.
7. Click OK.

Enabling and Purging the Audit Log

When you select the Enable Audit Generation option, the Server logs audit events for projects in the server configuration database. For example, the log records when change requests are created, and when a file is added. The audit log entries can be viewed from a client by selecting the Audit tab in the upper pane.

This operation can be performed only on a server that is running.

To enable the audit log:

1 Do one of the following:

- From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
- From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.

3 Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.

4 Select the Audits tab.

5 Select the Enable Audit Generation check box.

6 (Optional) To automatically delete entries after a specified length of time:

a Select the Purge Audit Entries Older Than check box. (Clearing this check box keeps the entries indefinitely.)

b Enter a number of days in the Days text box. The range is from 7 to 1000 days.

For example, to delete entries when they become approximately one month old, enter 30 days in the Day text box. When the server configuration starts, entries that exceed this purge limit are deleted.

7 Click OK.

Using Pooled Connections

Pooled connections are the concurrent connections that the server opens to the database engine. The range available is usually database-specific, and may also depend on the number of licenses you have for that database. The default value is 2, but an initial setting of 10–15 is recommended. The Server starts with the number you specify, and makes additional connections as needed.

This operation can be performed only on a server that is running.

To change the number of pooled connections:

1 Do one of the following:

- From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
- From the client, Select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 From the list of servers, select the server configuration that you want to change. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.

3 Click the Configure Server shortcut or Tools > Administration > Configure Server. The *Configure Server* dialog appears.

4 Select the Database tab.

5 Enter the new number of pooled connections in the Pooled Connections text box.

6 Click OK.

Verifying the Database at Startup (Native-I Vault Only)

When you start a server configuration that contains files in Native-I format, you can verify the database. This option does not exist for servers that have all of their files in a Native-II vault.

During the database verification process, the Server compares the tip revision (most recent revision) of each file in the archive with the information in the database about that revision, and logs any problems in the server's log file.

We recommend verifying the database only when you restore a backup of that database or when you are experiencing database problems. You should *not* perform this operation at every startup, because this can be a very time-consuming process, depending on the size of your database. If you are running the Server as a Windows service, you might never choose to verify the database at startup.

Tip If you use a Native-I vault, you can also use the Verify Native-I vault option in the Server Administration utility for verification purposes (when the server configuration is not running). Because the Verify Native-I Vault option is interactive utility, it can ask you to supply the location of a working file to substitute for a missing tip revision. The database verification process at startup is not interactive and cannot do this.

To perform tip revision verification for a server configuration with a Native-I format:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. From the list of servers, select the server configuration for which you want verification. The server configuration should not be running.

3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.

4. Select the Native-I Vault icon or Actions > Verify Native-I Vault. These options appear only if the server has a Native-I vault.

5. Select the Verify Database On Startup check box, to verify the database as part of starting the server configuration.

6. Select the Attempt Automatic Recovery check box, to make limited repairs.

   The verification process attempts to reconcile the differences between the server configuration's database and its archive, by doing the following:

   - If the tip revision described in the database is missing from the archive but can be found in the cache, the process adds it to the archive.
   - If the archive contains more revisions of a file than the database is aware of, the extra revisions are deleted.
   - The process backs up the archive before it deletes the revisions. Backups of the archive are in the archive folder, initially *RepositoryPath*\Vault\Archive. Backups of the cache are in the cache folder, initially *RepositoryPath*\Vault\Cache. The backup files are numbered and end in BU. The numbers for the archive and cache backups for a given verification attempt match.
   - If the revision information in the database differs from the actual cache contents, the extra files in the cache are deleted.

   Changes made during auto-recovery attempts are logged in the server's log file, which is stored in the folder you selected as the server configuration's repository.

7. Select the Stop Server On Errors check box, to prevent the server configuration from starting if the verification process reports any errors.

8. Click OK.

Changing Maximum Disk Cache Size and Cache Cleanup Interval (Native-I Vault Only)

In a Native-I vault, the disk cache folder contains copies of both the most recently checked-in and the most recently checked-out revisions of files. If a user checks out a revision that is already in the cache, the file is copied from the cache rather than from the archive. Using a disk cache speeds up the check-out process. The maximum cache size determines how many files can be stored.

By default, this folder is the Cache folder and is a child of the folder specified for the vault.

This operation can be performed only on a server that is running.

To change the maximum cache size:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so.

3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.

4. Select the Native-I Vault tab.

5. Enter the number of megabytes in the Maximum Cache Size text box. The default is 100 MB.

6. Enter a number of seconds between cache cleanup operations in the Cache Cleanup Interval text box.

   The cache is checked at this interval, and if the cache exceeds its effective size limit, the Server removes the least recently used files from the cache until the cache reaches 80% of that size limit. The Server reduces the cache size limit when there is insufficient disk space for the set size limit. The reduced size becomes its effective size limit. For example, suppose the current cache size is 60 MB, that the limit is 100 MB, and that only 20 MB of disk space is available. Then the effective limit to the cache size is 80 MB.

   The value for the cache cleanup interval may depend on the size limit of the cache, but we recommend using the default.

7. Click OK.

Setting the Vault File Lock Time Limit

When a user checks a working file into a Native-I vault, the Server locks that file's counterpart in the archive folder (by default a child of the vault folder). This ensures that no one else accesses that file until the check-in process is complete. If the check-in process fails, the file lock may not be automatically released. The time limit option for vault file locks enables you to specify an amount of time after which the Server releases any internal vault lock. The time limit should be long enough so that it does not force a normal check-in process to stop before the check in has completed. By default, the time limit is set to 60 seconds.

This operation can be performed on a server that is running.

To change the vault file lock time limit (Native-I vault only):

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

This action displays the *Server Administration* dialog.

2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the Native-I Vault tab.
5. Enter the number of seconds in the Vault File Lock Time Limit text box.
6. Click OK.

Changing Endpoints

The default TCP/IP port (endpoint) is 49201, but you can specify a different port for a server configuration. If you have more than one server configuration running on the same computer, each server configuration must use a unique endpoint. For example, if Server Configuration 1 uses the endpoint 49201, Server Configuration 2 must use a different endpoint.

If you attempt to run server configurations that have the same endpoint and computer name at the same time, only the first server configuration you select will start successfully. The remaining server configuration will appear to start, but in fact is ignored by the Server.

This operation can be performed only on servers that are running.

To change the endpoint for a server configuration:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the Protocol tab.
5. To activate a different port, enter its number in the TCP/IP Endpoint text box. The range for port numbers is 1023 through 65535.
6. (Optional) Click Default if you wish to return to the default endpoint setting (49201).
7. Click OK.

Setting Minimum Encryption for TCP/IP

Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines—such as the Internet. For TCP/IP connections, you can set a *minimum* level of encryption for a server configuration for IP addresses that access that server configuration. You can set different encryption levels for an IP address, ranges of IP addresses, or all IP addresses.

The encryption level can be set from a client on a per-workstation basis. Users must use at least the minimum level of encryption set for that server configuration.

This operation can be performed only on servers that are running.

To set an encryption level for transferred data, regardless of the IP address:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the Protocol tab.
5. In the TCP/IP Encryption Levels list box, select Default.
6. Click Modify. The *Set Default Encryption Type* dialog appears.
7. Select the type of encryption you want to use with the server configuration for IP addresses not specified in this list.
8. Click OK.

To set a different encryption level for a specific address or range of addresses:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the Protocol tab.
5. Click Add. The *Set Default Encryption Type* dialog appears.
6. Enter the starting IP address in the Starting IP boxes.
7. Enter the ending IP address in the Ending IP boxes.
8. Select the type of encryption to be used with the server configuration for these addresses.
9. Click OK.

Managing Event Handlers

StarTeamMPX has an event transmitter that must be installed on the same computer as the Server. In addition, StarTeamMPX Services (Message Broker and Multicast Service) can be installed on the same or other computers, depending on your needs.

When you install the event transmitter, the following events appear in the *Event Handlers* dialog:

- Unicast On-site, if you have installed the Message Broker
- Multicast On-site, if you have installed the Multicast Service (which can be used only by the Windows client)

For more information about this server, its XML files, properties, and values, see the *StarTeamMPX Administrator's Guide*. This section explains how to add, modify, and delete event handlers. It does not explain the purpose of the properties, or the range of values that can be assigned to them.

The *Event Handlers* dialog provides a simple interface for editing the StarTeamMPXTransmitter.XML files.

This operation can be performed only on servers that are running

To assign default event handlers for the server and/or clients:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the Event Handlers tab.
5. Select an existing event handler.
6. Do one or both of the following:
   - Click Server Default, to make the selected profile the profile for the server. A server icon appears in front of the default server profile.
   - Click Client Default, to make the selected profile the default profile for clients. A green check mark appears in front of the default client profile.

     As users create server descriptions on their workstations, the profile selection defaults initially to this profile. Users can change from the default to another existing profile.

   If a profile is both the server and client default, you see only the server icon.
7. Click OK.

To review or modify an existing event handler:

1. From the Event Handlers tab of the *Server Configuration* dialog, select an existing event handler.
2. Click Modify. The *Event Handler Profile Properties* dialog appears. This dialog allows you to review the property settings, change the settings and add or remove properties.
3. To change a setting:
   a. Select a setting from the Profile Properties list box.
   b. Click Modify. The *Event Handler Property* dialog appears.
   c. Change the value.
   d. Click OK.
4. To add a property:

a. Click Add. An empty *Event Handler Property* dialog appears.

b. Enter a property name and value in the appropriate check boxes.

c. Click OK.

5. To remove a property:

a. Select a setting from the Profile Properties list box.

b. Click Remove. Be aware that you cannot delete a profile that is currently used as the default profile.

6. When you are finished, click OK.

To remove an event handler:

1. From the Event Handlers tab of the *Server Configuration* dialog, select an existing event handler.
2. Click Remove.

To create new event handlers:

1. From the Event Handlers tab of the *Server Configuration* dialog, select an existing event handler.
2. To create a new event handler from scratch:

a. Click Add. An empty *Event Handler Profile Properties* dialog appears.

b. Enter a name and description in the appropriate text boxes.

c. Click Add to display an empty *Event Handler Property* dialog.

d. Enter the property name and its value in the text boxes.

e. Repeat steps c and d until you have added all the properties you need.

3. To create a new event handler from an existing one:

a. Select an existing event handler that is very similar in its properties to the new handler that you need.

b. Click Copy. The *Event Handler Profile Properties* dialog displays the properties for the selected event handler.

c. Change the name and description in the appropriate text boxes.

d. Select and modify other properties as appropriate.

4. When you are finished, click OK.

Enabling Directory Service Support

StarTeam now allows password verification with Microsoft Active Directory. Active Directory service is included with Microsoft Windows Server 2003 and Microsoft Windows 2000

Server operating systems. It allows centralized, secure management of an entire network.

To validate users against the directory server, the Server must be on a trusted domain in relation to that server. On the Directory Service tab of the *Server Configuration* dialog, you must also select the **Enable directory service** option and enter the location and port number of the directory server.

For each individual who will be validated against the directory server, you must select the **Validate with directory service** option on the *New User Properties* or *User Properties* dialog and enter a Distinguished Name (used to uniquely identify a directory services user). See “Adding Users” on page 55 for more information.

Even if the settings are correct, the user will not be able to log on if the directory server is unavailable. Although directory service support is off by default, it can be activated at any time. The server cannot be running at the time you enable or disable the support.

When the user supplies a StarTeam logon name and a Microsoft Active Directory password, StarTeam Server recognizes that the user is set up for directory service password validation and uses the **Distinguished Name** and password as it contacts Active Directory. If the password is verified, the user is allowed to access the StarTeam Server.

To enable directory service support:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. Select the Directory Service tab.
5. Select the **Enable directory service** option. By default this option is not selected.
6. Enter the Host name and SSL port of the directory server. Both must be specified.
7. Enter a secure port number.
8. Click OK. The system displays a message instructing you to reboot the server. This action must be performed for directory service support to be enabled.

Note Remember that a user cannot be authenticated by the directory server unless the Validate through directory service option is selected on the Logon tab of the *New User Properties* or *User Properties* dialog and a **Distinguished Name** is entered for that user. See *"Adding Users" on page 55* for more details.

Diagnosing Server Problems

To reduce the amount of time spent diagnosing problems, the application provides tracing and debugging tools for the server. It can create either, or both, trace command files and diagnostic (.dmp) files.

Both of these options are turned off, by default. If you encounter a problem, you can simply turn them on and create files that you can review or discuss with Borland support.

- Trace Commands

  The trace option create a file that records single server commands. Commands to be traced must have a duration time that equals or exceeds the number of specified milliseconds. The default time is 0. If you wish to record only commands of longer duration, you should adjust this setting, to avoid taking up unnecessary space in the trace file.

  No trace file should generate more than 10 MB of data per day. Typically, customers see a only small fraction of this amount of data per day.

  Trace data is stored in a Server.trc file, which consists of a header followed by an arbitrary number of records. When a trace ends, the server timestamps the existing

file as Server.*time*.trc. Trace files are located in the *repositoryPath*\Log\Trace folder. The next trace file starts when the server configuration is restarted or the trace option is turned on.

- Diagnostic (.dmp) Files

  The application creates some minidump files automatically, while others are created only when the.dmp options are turned on. Minidump files can be created for either or both:

  - Asserts (unexpected conditions). Server log entries with code number 8.
  - Exceptions (errors, typically access violations). Server log entries with code number 4.

  Minidump files are created in the same location as the server log file. The general naming convention for these files is *prefix-counter-time*.dmp, in which *prefix* identifies the source of the dump, *counter* is an integer that increments with each .dmp file to ensure that names are unique, and *time* identifies the local server time at which the dump was created.

This operation can be performed only when the server is running.

To turn on the options used for server problem diagnosis:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server. The *Configure Server* dialog appears.
4. To create .trc files:
   - Select the Trace operations that take at least __ milliseconds option.
   - If you do not want to use the default of 0 milliseconds, enter a different number.

   By default, this option is not selected.
5. To create .dmp files, select either or both of the following options:
   - Unexpected conditions (server log entries with code #8)
   - Errors (server log entries with code #4).

   By default, these options are not selected.
6. Click OK.

Changing Server Session Options

The session options for each server configuration are stored in the starteam-server-configs.xml file. You can modify a number of these options from the Server Administration utility or the command prompt with the starteamserver command. For information about this command, see Appendix D, “Using the starteamserver command,” on page 253.

When the server configuration is not running, you can modify the following session options with the Server Administration utility:

- Server configuration name
- Log file path
- Database connection information
  - DSN name
  - User name
  - Password

Any changes you make will take effect the *next* time you start the server configuration. You can also change certain configuration options by using the Start With Override button. For more information, see "Starting a Server Configuration" on page 15.

To change the session options for a server configuration:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you are using the client, you will be able to administer remote servers only.
   a. If the server configuration is running, click the Shut Down Server icon or Actions > Shut Down Server on the menu.
   b. Click the Properties icon or Server > Server Properties on the menu.

   The *<server configuration> Properties* dialog appears.
3. To change the server configuration name, enter a new name in the Configuration Name text box.
4. To change the log file path:
   a. Click Change Path.
   b. Select a new folder for the server log file (Server.*locale*.Log).
   c. Click OK.
5. If you are changing the database or schema user used by the server configuration, you can change the database connection information:
   a. Select the Database Connection Information tab.
   b. Enter a new DSN name in the Server Name/DSN text box.
   c. Enter a new user name and password in the User Name and User Password text boxes.

      If the server configuration uses an Oracle database, these boxes are named Schema User Name and Schema Password.
   d. Click Verify Connection to be sure that the DSN name, user name, and password connect to the database.
6. Click OK.
7. Restart the server configuration to see the changes take effect.

Reviewing Database Information

You can review the database type and DSN name from the *<server configuration> Properties* dialog or the *Server Configuration* dialog. The *<server configuration> Properties* dialog also shows the user name, provides data file and transaction log information, and enables you to verify the connection to the database.

To see server configuration database information from the *Properties* dialog:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Properties icon or select Server > Properties from the menu. The *Server Properties* dialog appears.
4. Select the Database Connection Information tab. This action displays the database name, DSN, and user name. You can click Verify Connection to check that the DSN name, user name, and password correctly connect to the database.
5. (Only for server configurations using MSDE or SQL Server databases) Select the Data Files and Transaction Logs tab.

   You will see the size and location information for data and transaction log files in the database used by this server configuration.
6. (Only for server configurations using Oracle databases) Select the Tablespace Information tab.
7. After you have finished viewing the information, click OK.

To see server configuration database information from the *Configure Server* dialog:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. From the list of servers, select the server configuration that you want to modify. The server must be running.

   If you have not yet logged on, you will be asked to do so. If you are using the client, you will be able to administer remote servers only.
3. Click the Configure Server shortcut or select Tools > Administration > Configure Server from the menu. The *Configure Server* dialog appears.
4. Select the Database tab. This tab displays the database name, DSN, and number of pooled connections.

Note Older StarTeam Servers may have Access databases that are no longer supported for upgrades.

Chapter 4

Working with the Application Logs

The application's clients and servers generate a number of log files. These logs enable an administrator to evaluate the performance of the system and potentially troubleshoot problems.

Each server configuration has its own server log and security event log. Each client creates its own .Log file, which records activity between that client and the server configurations it is connected to. This chapter explains how to display and work with each of these logs.

Users must have the appropriate security access rights in order to view a log file. These access rights can be set using the Access Rights button on the *Server Administration* dialog. For more information about access rights, see Chapter 12, "Controlling Access to Objects," on page 189.

Reviewing the Server Log

The server log file (Server.*locale*.Log) records the activity on a server configuration. Each time you start a server configuration, the Server renames the existing log file and creates a new log file for the current server configuration session. The log file from the previous startup is renamed to include the date and time at which it was renamed (Server.*locale.date*.Log). For example, if you start a server configuration on November 9, 2005 at 5:22 P.M., the old Server.*locale*.Log file is renamed Server.en-US.2005-11-09-17-22-59.Log, and a new Server.*locale*.Log file is created whose time stamp might be 11/9/2005 17:23:03.

If the locale specified for the operating system on which your server runs is not US English, you will have two server log files: one for US English and one for your locale. For example, you might have both Server.en-US.Log and Server.fr-FR.Log. The first log is for support purposes, and the second log is for your use.

You can view the contents of the server log file at any time, even while the server configuration is running. Only the last 64K of the log file appears. To see the entire file, use Notepad, WordPad, or another text editor to display it.

To determine the location of a server log file from the Server Administration utility

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.
3. Click the Configure Server shortcut or Tools > Administration > Configure Server. These actions display the *Configure Server* dialog.
4. Look at the top of the General tab to find the location of the log file.

To review the contents of a server log file:

1. Do one of the following:
   - From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions displays the *Server Administration* dialog.
2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.
3. Click the Server Log shortcut or Tools > Administration > Server Log. These actions display the Server Log contents.

   The format provides a line number, code, date and time, and message. The code numbers are arbitrarily assigned and not necessarily in order of severity. They represent the following:

   00000001 Message
   00000002 Warning
   00000004 Error
   00000008 Unexpected Condition
4. To display only the errors in the log, select the Errors Only check box.
5. To review the most recent entries, click Reload.

Copying and Printing Data from the Server Log

On Windows systems, you can copy data from the log window to the Windows clipboard. From the clipboard, you can paste the data into other applications, such as Microsoft Word or Notepad.

To copy data from the server log:

1. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane and select the Server Log Shortcut.
   - From the menu, select Tools > Administration > Server Log.
2. Select the data that you want to copy.
3. Press *Cntrl + C*.

4 Click OK to exit the dialog.

5 Using Cntrl + V, paste the information into a text editor or word processing application.

6 Click File > Print from the menu to print the data.

Notifying the System Managers Group

Any users in the System Managers group, a default, will receive e-mail when an error is logged in the server log. This group is initially empty. To add users to the System Managers group, follow the steps in “Adding Users” on page 55.

For information on changing the language of messages sent to the System Managers group, see “Enabling E-mail Notification” on page 29.

Reviewing the Security Event Log

If you have access rights to a server configuration, you can view its security event log at any time. The security event log is not a typical .Log file, as its data is stored in the application database. This operation can be performed only when the server is running.

To view the security event log:

1 Do one of the following:

- From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.
- From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote configurations only.

3 Do one of the following:

- Click the Accounts bar in the lower left pane, then select the Security Log shortcut.
- Select Tools > Accounts > Security Log. These actions display the Security Log contents.

This log lists each secured event (such as logging on or off), the date and time it occurred, the user performing the operation, the workstation from which the operation was performed, the item acted upon, and whether the operation failed.

4 Use the Security Event Type drop-down list box to view all events of a particular type.

Table 4.1, “Security Event Types” presents a list of security event types and a description of these events.

Table 4.1 Security Event Types

Event Type	Description
Acquire item ownership	Ownership is not implemented at this time.
Add item owner	Indicates that a user created a folder or an item.
Add user/group	Indicates that a user or group was added to the server configuration.

Table 4.1 Security Event Types (continued)

Event Type	Description
Add/Edit container access rights	Indicates that access rights were added or changed for a group of objects contained in another object. For example, if you select Project > Access Rights and change rights for all change requests in the project, that event fits into this category.
Add/Edit item access rights	Indicates that access rights were added or changed for a specific object. For example, if you change access rights for a project, that event fits into this category.
Change user	Indicates that someone changed user names as part of a replication process. This event can occur when special clients, such as Notification Agent, perform operations.
Delete container access rights	Indicates that access rights were deleted at the container level.
Delete item access rights	Indicates that access rights were deleted at the item level.
Delete item owner	Ownership is not implemented at this time.
Delete user/group	Indicates that a user or group was deleted.
Edit user/group	Indicates that the properties for a user or group were changed in some way.
Force user logoff	Indicates that a user was forced to log off the server configuration.
Item access check	Indicates that access rights were checked to see if the user could access a specific item.
Logoff	Indicates that a user logged off the server configuration.
Logon	Indicates that a user logged on to the server configuration.
Logon attempt: Account lockout	Indicates that a user attempted to log on and the account was locked.
Logon attempt: Expired password	Indicates that a user attempted to log on and the password had expired.
Logon attempt: No such user name	Indicates that a user attempted to log on with a non-existent user name.
Logon attempt: Restricted access time	Indicates that a user attempted to log on at a time when he or she was not allowed access.
Logon attempt: Suspended account	Indicates that a user attempted to log on and the account was suspended.
Logon failure	Indicates that an incorrect password was used during the logon process.
Policy change	Indicates that a system policy has changed.
User status change	Indicates that an administrator suspended, reactivated, locked, unlocked, or required a password change on a user's account.

To reload the security event log and review the most recent entries:

- From the *Security Event Log* dialog, click Reload.

To print the data selected from the log:

- From the *Security Event Log* dialog, click Print Selection.

Purging Entries from the Security Event Log

Depending upon the number of users and the amount of activity on a server configuration, the security event log may grow rapidly. To keep the log to a reasonable size, you can have the Server delete old entries. First, decide how long you want to have security events available, then configure the server configuration to purge entries that are older than this time period.

This operation can be performed only when the server is running.

To set the interval for purging the security event log:

1. From the computer that has the Server installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server.

   This action displays the *Server Administration* dialog.

2. Select the appropriate server configuration from the list of servers. If you have not yet logged on, you will be asked to do so.

3. Do one of the following:

   - Click the Accounts bar in the lower left pane, then select the System Policy shortcut.
   - Click Tools > Accounts > System Policy from the menu.

   The *System Policy* dialog appears.

4. Select the Security Events tab.

5. Select the Purge Security Event Entries Older Than __ Days check box. (Clearing this check box keeps the entries indefinitely.)

6. Enter the number of days in the text box. The range is 30 to 1000. The default is 180.

   The next time the server configuration starts, entries that exceed the purge limit are deleted.

7. Click OK.

8. Restart the server configuration for the purge interval to take effect.

Reviewing the StarTeam.Log File

The StarTeam.Log file records the operations performed on your Windows client workstation during a session. It helps you troubleshoot and document errors or operations between the server and your workstation that failed during server configuration sessions. The file may contain the following types of information:

- Commands sent by your workstation to a server configuration when you open and work with a project. If you work with projects on several different server configurations, you can configure the StarTeam.Log file to include the server configuration name with the command information it records.
- Commands performed locally on your workstation, such as setting personal options.
- Error messages generated while using the application.
- Events performed by StarTeamMPX.

Every time you start your client, the system creates a StarTeam.Log file in the folder location specified in your personal options.

On most systems, the default location for the StarTeam.Log file is C:\Program Files\Borland\StarTeam *x.x*. If there is a StarTeam.Log file already in this folder, The application renames the existing file to include the date and time at which it was renamed. For example, if you create a StarTeam.Log file on November 9, 2005 at 10:35 A.M., the old StarTeam.Log file is renamed StarTeam-09-Nov-05-10-35-18.Log, and a new StarTeam.Log file is created.

Tip Because the application creates a new StarTeam.Log file every time you start the client, the log folder can fill up quickly.

To control the number of log files in the folder, you may want to periodically delete old log files from the output folder or disable the StarTeam.Log option. To disable the option, clear the Log Errors and the Log Operations check boxes on the Workspace tab of the *Personal Options* dialog

To display the StarTeam.Log file (Windows client only):

- From the client, select Tools > StarTeam Log File from the menu bar.

You can also import and view the data from a StarTeam.Log file using any application that supports tab-delimited fields. For example, if you save the file with a .csv extension, the file can be opened in Microsoft Excel.

Customizing the StarTeam.Log File

The Workspace tab on the *Personal Options* dialog enables you to specify the location and the type of data recorded in the StarTeam.Log file (Windows client only).

To customize the StarTeam.Log file:

1. From a client, select Tools > Personal Options. The *Personal Options* dialog appears.
2. On the Workspace tab, enter or browse for the location of the StarTeam.Log file in the Log Output Path text box. The default is the location in which the application is installed; for example, C:\Program Files\Borland\StarTeam client_name\Log

   The current log file is always named StarTeam.log. Log files from earlier sessions of the application include the date and time the file was last modified.

Note: StarTeam.log contains data about operations sent from your workstation to one or more servers, depending on what project views you have open. This data includes the name of the project so that you can isolate data for a particular server, when necessary.

3. Select the types of data you want to include in StarTeam.Log.
   - Log errors

     Set by default. Records errors that occur while you are using the client. The errors log lists the date and time you started your server configuration and any errors or failed operations between the server and client. The application identifies each failed operation by an internal ID and provides an explanation.

     For example, you might see:

     ...Operation 40956 failed: TCP/IP Socket Error 10054:...

   - Log operations that take at least ___ milliseconds

     Select this option to record operations that take longer than a specific number of milliseconds. (An operation is a command that results from user actions. Operations can be executed on either the Server or the client.) The milliseconds time setting stops the log from filling up with operations of little importance. The default is 10 milliseconds.

     This option records information on the date, time, and UI Operation number for each command executed by your workstation. You can record either Summary or Detail information.

     - Selecting Summary records the time spent for the total operation, client execution time, and server execution time. The application identifies each operation by an internal ID, such as Statistics for Operation 40001.
     - Selecting Detail records a detailed breakdown of all server commands performed for each operation. The log identifies the server address, project, and component (File, Change Request, Requirement, Task, or Topic) for each server command. The application identifies each server command by an internal ID, such as Public Server Command 10.

- Log StarTeamMPX events

  Selecting this option records information about StarTeamMPX events for this client. The log identifies the time and date on which a StarTeamMPX event (an automatic refresh or file status update) took place. The log prefaces a StarTeamMPX event as “Statistics for Events” and uses internal IDs and brief explanations to identify the server event.

  The following example describes a status change for a file:

  ...Statistics for Events /1b21dd1-e208-51ea-01b2-1dd1e20851ea/Object/File/Modify

  You can log StarTeamMPX events only if you check the “Enable StarTeamMPX” checkbox on the StarTeamMPX tab. For StarTeamMPX-related operations, any changes you make on the Workspace tab do not apply to projects already open. However, the application will log StarTeamMPX events for any projects you open from this point forward.

4 When you are finished, click OK.

Chapter

5

Setting Up Users and Groups

All of the application clients enable you to add new groups and users manually. When users log onto the application, they can now be validated by a password that has been entered in or imported to the application or obtained from Microsoft Active Directory Services (the LDAP server). This operation is possible only when the server is on a trusted domain in relation to the LDAP server.

Understanding the Default Groups

New server configurations come with a few default groups: All Users, Administrators, System Managers, Security Administrators, and Collaboration Users (for StarDisk). These groups come with default privileges but can be assigned privileges in accordance with your company policy.

The users in the Administrators group initially have all available privileges, giving them complete access to the system unless the system is set up to ignore privileges. The All Users, System Managers, Security Administrators, and Collaboration Users groups initially have no privileges. See “Adding Groups” on page 52 for more information about privileges.

- All Users

  All users are members of the All Users group because All Users is the root group in the User manager and because all members of a child group are members of its parent group.Therefore, all users inherit any rights and privileges assigned to this group.

- Administrators

  This group initially contains the Server Administrator user. You may want to add others who have administrative privileges.

- System Managers

  The users in this initially-empty group receive e-mail (at the address specified for them in the User Manager) whenever an error is added to the server log.

- Security Administrators

  The users in this group can receive e-mail about users who were locked out or attempted to log on unsuccessfully. For details, see “Setting the Number of Logon

Attempts” on page 71.This group initially contains only the user who has been designated as the Server Administrator.

- Collaboration Users (for StarDisk)

  The users in this initially-empty group can access the StarDisk project in a server configuration if that project exists and if the users have been invited to join that project. These users actually access only a subset of the StarDisk project.

Caution Never have only one user account with administrative privileges. If you are logged on using the only user account with administrative privileges and you become locked out, you have no way to unlock your own account.

The Server comes with a user named “Administrator” who has the password “Administrator.”. Because this is common knowledge, you will want to change at least that password.

Adding Groups

Users who can log onto a server configuration can be organized into groups. Creating and using groups simplifies the task of managing security on a project, because each group can be assigned a set of privileges that apply to all the users in that group, rather than setting privileges on a user-by-user basis.

The status bar on the *User Manager* dialog displays the number of users in the selected group who have access to the server configuration, the number of users connected to the server configuration, and the number of users logged on. The number of users connected to the server configuration and the number of logged on users differ when individual users log on more than once.

This operation can be performed only when the server is running.

To add a group:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
   - Select Tools > Accounts > User Manager from the menu.

   These actions display the *User Manager* dialog.
4. Select a group from the Groups tree to be the parent of the new group.

Note We recommend that, initially, you select the All Users group when adding a new group. Subsequent groups can be added to any group listed under the All Users group. Avoid adding new groups to the administrative and management group.

If a user is a member of a child group, it is also implicitly a member of the parent group—even if the member’s name does not appear in the list when you select the parent group. You must select the Show Users in All Descendant Groups check box to see the complete list of members for a selected group that has child groups.

5 Click New Group. The *New Group Properties* dialog appears.

6 Enter the group name in the Name text box.

7 Enter a description of the group in the Description text box.

8 Select the Privileges tab. The privileges selected on the Privileges tab can override any Access Rights that have been previously set for any user in the privileged group. However, the privileges are not a substitute for Access Rights. If you have not set up Access Rights, you have no security system.

The privileges set on the Privileges tab apply to all objects in all projects in a server configuration. For example, if you give a group the Delete Item privilege, any user in that group can delete any project, view, folder, child folder, or item from the server configuration, regardless of what the Access Rights are for deleting these items.

9 Set privileges as appropriate. For additional information about security, see "Setting Group Privileges" on page 193.

10 Click OK. The new group appears in the Groups list.

Changing the Parent of a Group

Groups are arranged hierarchically in the Groups tree. You can change a group's position in that hierarchy. This operation can be performed only when the server is running.

To change a group's position in the Groups tree:

1 D0 one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the group to be moved from the Groups tree.

5 Right-click and select Change Parent Group from the context menu. The *Change Parent Group* dialog appears.

6 Select a new parent group, then click OK.

7 Click OK.

Determining Who Is a Member of a Group

A group has both explicit and implicit members. Explicit members are those users who were specifically assigned to the group. Implicit members are users who belong to the group because they are explicit members of one of the child groups included in that group. This operation can be performed only when the server is running.

To determine the members of a group:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the group from the Groups tree. The explicit members of the group appear in the Users list box.

5 Select the Show Users in All Descendant Groups check box to also display the implicit members of the group in the Users list box.

Tip If you want to determine the groups to which a particular user belongs, see "Changing a User's Group Membership" on page 63.

Removing an Empty Group

Removing a group does not remove any user accounts from the server configuration; it only removes the group name from the list of available groups. Prior to removing a group, you must either delete all users from this group, or move the users to another group, because the application will only delete *empty* groups.

This operation can be performed only when the server is running.

To remove a group:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select a group from the Groups tree.

5 Right-click to open the context menu, and select Delete.The system displays the following message: Do you want to delete group groupname?

6 Click Yes.

- If the group is empty, it is removed from the Groups list box.
- If the group contains users, the system displays the following message:

  ```
  The group you want to delete contains user accounts. Please delete these
  user accounts or move them to another group prior to deleting a group.
  ```

  When you see this message, click OK. Then either delete the users in this group or move them to another group.

Adding Users

If you have the appropriate access rights, you can add users to a server configuration from either the Server Administration utility or a client. Initially, you add a user to a specific group, such as Developers or Testers. The user becomes an explicit member of this group and an implicit member of any of this group's parent groups, such as the All Users group. For more information about explicit and implicit membership, see "Changing a User's Group Membership" on page 63.

This operation can be performed only when the server is running.

Caution Creating a user account with the name "StarTeam" has been known to cause problems when using the command line stcmd server-mode command to lock or unlock the server configuration. The command requests a password even when the user has a blank password or when a password has already been provided.

To add a user:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select a group from the Groups tree.

5 Click New User. The *New User Properties* dialog appears.

6 Enter the user's name in the Full Name text box.

7 (Optional) Enter the user's e-mail address in the E-Mail text box using one of the following formats:

```
yourname@company.com
```

OR

```
user_friendly_name<yourname@company.com>
```

In the Windows client, if the E-mail text box blank in the *Server Administration* dialog has been left blank, the following message displays after you click OK.

```
Attempt to resolve e-mail address?
```

- Click Yes to have the application search your address server for an e-mail address.

  If the application cannot locate an e-mail address for this user, the following message displays:

  ```
  E-mail address cannot be resolved. The user will be excluded from e-mail
  related operations.
  ```

  If the application is unable to locate an e-mail address for the user, you will have to manually enter it. Otherwise, the user will be excluded from operations that use e-mail.

- Click No if you do not want the user to have an e-mail address.

Note If the e-mail address listed for a user is ambiguous or incorrect, your mail application will consider e-mail to that user to be undeliverable. For example, if two people have e-mail addresses that begin with “jmarsh”, or if the user “jmarsh” has two e-mail addresses known to the mail application, this problem cannot be resolved uniquely.

8. (Optional) Enter the user’s phone number, voice mail number, pager number, fax phone number, and street address in the appropriate text boxes.
9. Select the Logon tab.
   a. In the User Name text box, enter the name to be used to log onto the application.

      If you enter a user name that already exists, the following message displays after you click OK:

      ```
      A user with a given user name already exists.
      ```

   b. If you want to validate the user against the server, select the Validate through StarTeam Server button.

      Enter a StarTeam password for the user in the Password text box and again in the Confirm text box. Asterisks appear in the text box instead of the password itself.

      - If the password’s minimum length can be zero, you do not have to enter a password.
      - If you are using strong passwords, be sure to follow the rules for those passwords. See page 69 for more information.

   c. (For Microsoft Active Directory or OpenLDAP) To validate the user against your organization’s directory server, select the Validate through directory service button. Then enter the Distinguished Name for the user. An alphanumeric value of up to 254 characters, this value is used to uniquely identify the directory services user.

Note To use directory service validation, the Server must be on a trusted domain in relation to the LDAP server.

10. (Optional) Select the Access Policy tab and specify when this user can access the server configuration.
    a. Select one of the following option buttons:

       - Access not restricted (the user can access the server configuration at any time)
       - Standard five day work week (the user can access the server configuration Monday through Friday from 8 A.M. to 5 P.M.)

- Custom access hours (to set one or two time periods per day when the user can or cannot access the server configuration)
- Select a day of the week from the Day list box.
- Select the No Access on That Day check box to deny access, or clear it to allow access on that day.
- Use the From and To boxes to set one or two time periods when access is either allowed or denied.

b. If the user's workstation is not in the same time zone as the computer on which the server configuration is running, select the Adjust for Workstation Time Zone check box, and enter the number of hours from Greenwich Mean Time (GMT) in the "hours from GMT" field.

11. Click OK.

12. Add the new user explicitly to other groups, as appropriate.

    Remember that a user is implicitly already a member of the current group's parent groups, but you must explicitly add a user to groups that are not parents of the current group. See "Changing a User's Group Membership" on page 63 for more information.

Licensing

You can use Borland License Server or FLEXlm with StarTeam Server. This gives you a choice between using a license server and using the StarTeam licensing found in this and earlier releases.

If you use a license server, users must use their network logon names as their StarTeam user names.

As the Administrator, you will:

- receive licensing information from Borland via email (a sales representative should put this in motion)
- configure the server (this is covered in the license server documentation. The link is in the License Certificate email)
  - install the license server and host the licenses sent by Borland (this involves accessing a Borland web site and downloading Borland license files called slips)
  - place the slips in the /license folder, a subdirectory of the StarTeam Server installation folder.
- configure the license server for users (this is covered in the license server documentation)
- use the StarTeam Server Administration tool to:
  - change user names to network logon names.
  - assign users to specific licenses

When StarTeam Server starts up, it checks for slips and stores information about them in memory. It does not recognize new slips until the next restart.

When a user logs in from a StarTeam client application, StarTeam Server tells the client application what features are available to its user based on the license assigned to that user.

If the user is assigned a license from a slip, but that slip is no longer in the /license folder, StarTeam Server displays an error message. If the user's license type is Unassigned, the user is not logged on and StarTeam Server returns an exception.

Licensing Users Using Borland License Server

Once you have downloaded all of the .zip files containing the .slip files per the instructions provided on the web, copy each concurrent_.slip and/or named_.slp file into your License directory (C:\Program Files\Borland\StarTeam Server\License). Copy each server_.slip file to(C:\Borland\BLS4\conf).

The following procedures explain how to assign licenses to specific users. Once the slips have been placed in the /license folder, User Manager can display information about the slips and an administrator can assign licenses from those slips to users.

To assign licenses to existing users:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
   - Select Tools > Accounts > User Manager from the menu.

   These actions display the *User Manager* dialog.
4. Select one or more users.
5. Select User > Assign License to display a list of the licenses that can be assigned.

The license types are:

- StarTeam named user licenses. This is the same as named user licenses in earlier StarTeam releases.
- StarTeam concurrent user licenses. This is the same as concurrent user licenses in earlier StarTeam releases.
- Unassigned. Select this “license type” when a user has no license.

Notes

The status line at the bottom of the User Manager window provides licensing statistics.

You may need to edit the properties for the user to change the user name to the network logon name.

To assign a license to a new user:

1. From User Manager, select the New User Button. The *New User Properties* dialog opens.
2. On the General tab, select a license type from the License drop-down list box. The default is StarTeam Concurrent.
3. Enter the rest of the data on the General and other tabs as appropriate. Remember to use the network logon name for the User Name text box on the Logon tab.
4. Click OK to exit the *New User Properties* dialog box.

Licensing Users Using StarTeam

Licenses determine how many users can access the Server. Users can have either named user or concurrent licenses. A named user license (formerly called a fixed license) can be used only by the user who has been assigned that license. For example, if you have 5 named user licenses and 25 concurrent licenses, the five users who receive the named user licenses are guaranteed access to the server. No one else can use their licenses.

A concurrent license (formerly called a floating license) can be used by any user who does not have a named user license. For example, users without named user licenses receive concurrent licenses on a first-come, first-served basis. After all the concurrent licenses are in use, users attempting to log on are notified that there are no more licenses available at this time. They can try again later.

When you first register your server, you enter one or two serial numbers: one for named user licenses, and/or one for concurrent licenses. The licenses also indicate what components and features the server can provide for users. When you use multiple serial numbers, they must all identify the same edition of the application. The Standard edition supports named user licenses only.

Starting with the 5.4 release of StarTeam Server, you can add more named user or concurrent licenses by entering additional serial numbers. The Server keeps track of the total number by summing the licenses supplied in each serial number. This is referred to as *stackable* licensing.

Previously, to add named users, you had to delete your old named user license serial number and add a new one with the total number of named users. You also had to replace your old concurrent license serial number with a new one indicating the total number of concurrent licenses. Therefore, you never had more than two serial numbers at one time.

If you have an evaluation license, you must delete it before entering the serial numbers for either named user or concurrent licenses. If you are upgrading from a previous release, you must delete the serial numbers for that release.

You can add as many users as you choose, but access to the server is granted only to users with named user licenses or to users who receive concurrent licenses as they log on. If you have named user licenses, you must assign them to specific users in the *User Manager* dialog. An anchor appears before the name of users with named user licenses. Before assigning named licenses, you must add the users. See “Adding Users” on page 55.“

The Server Administrator is automatically assigned a named user license which cannot be removed. This free license is *not* counted against the number of named user licenses you have available.

After the server is licensed, named-user licenses can be assigned.

Tip The *User Manager* dialog’s status bar indicates how many named and how many concurrent licenses are in use.

To assign a named user license:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Locate the user to whom a named user license will be assigned. To do this, you may need to select a particular group, or select the Show Users in All Descendant Groups check box.

5 Right-click, and select Add Named User License from the context menu.

After a named user license is assigned to a user, an anchor appears before the user's name.

To remove a named user license:

1 From the *User Manager* dialog:

a Select one or more users.

b Right-click, and select Remove Named User License from the context menu.

Assigning Licenses to Specific Users

The following procedures explain how to assign licenses to specific users. Once the slips have been placed in the /license folder, User Manager displays information about the slips and an administrator can assign licenses from those slips to users.

To assign licenses to existing users:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 From the *User Manager* dialog:

a Select one or more users.

b Right-click and select Assign License from the context menu to display a list of the licenses that can be assigned.

The license types are:

- Licenses from slips. The figure above shows slip 9954 which contain named user licenses, but slips can contain either named or concurrent user licenses.
- StarTeam named user licenses. This is the same as named user licenses in earlier StarTeam releases.

- StarTeam concurrent user licenses. This is the same as concurrent user licenses in earlier StarTeam releases.
- Unassigned. Select this "license type" when a user has no license.

Notes The status line at the bottom of the User Manager window provides licensing statistics.

You may need to edit the properties for the user to change the user name to the network logon name.

To assign a license to a new user:

1. From the *User Manager* dialog, select New User. The *New User Properties* dialog opens.
2. On the General tab, select a license type from the License drop-down list box. The default is StarTeam Concurrent.
3. Fill in the rest of the data on the General and other tabs as appropriate. Remember to use the network logon name for the User Name text box on the Logon tab.
4. Click OK to exit the *New User Properties* dialog.

Changing a User's Password

In addition to setting or changing a user's password, you can specify how long a password is usable, how many characters a password must have, and whether strong passwords are required. See "Forcing a Password Change" on page 69.

This operation can be performed only when the server is running.

To change a password:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
   - Select Tools > Accounts > User Manager from the menu.

   These actions display the *User Manager* dialog.
4. Select the user's name. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:
   a. From the Groups tree, select the All Users group.
   b. Select the Show Users in the All Descendant Groups check box.
5. Right-click, and select Properties from the context menu. The *User Properties* dialog appears.
6. Select the Logon tab.
7. Verify that the Validate through StarTeam button has been selected.
8. Enter a new StarTeam password for the user in the Password text box.

9 Enter the password again in the Confirm text box.

10 Click OK.

The new password will become effective the next time the user logs on.

License Server Support

You can now use Borland License Server or FLEXlm with StarTeam Server. This gives you a choice between using a license server and using the StarTeam licensing found in this and earlier releases.

If you use a license server, users must use their network logon names as their StarTeam user names.

An administrator performs the following steps.

1 Receives licensing information from Borland via email (a sales representative should put this in motion)

2 Installs the license server (the license server documentation explains how to do this)

3 Hosts the licenses sent by Borland (this involves accessing a Borland web site and downloading Borland license files called slips)

4 Places the slips in the /license folder, a subdirectory of the StarTeam Server installation folder.

5 Configures the license server for users (this is covered in the license server documentation)

6 Uses the StarTeam Server Administration tool to:

   a Change user names to network logon names.

   b Assign users to specific licenses

When StarTeam Server starts up, it checks for slips and stores information about them in memory. It does not recognize new slips until the next restart.

When a user logs in from a StarTeam client application, StarTeam Server tells the client application what features are available to its user based on the license assigned to that user.

If the user is assigned a license from a slip, but that slip is no longer in the /license folder, StarTeam Server displays an error message. If the user's license type is Unassigned, the user is not logged on and StarTeam Server returns an exception.

The following procedures explain how to assign licenses to specific users. Once the slips have been placed in the /license folder, User Manager can display information about the slips and an administrator can assign licenses from those slips to users.

To assign licenses to existing users:

1 Do one of the following:

   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 From the *User Manager* dialog:

a Select one or more users.

b Right-click, and select Assign License from the context menu to display a list of the licenses that can be assigned.

The license types are:

- Licenses from slips. The figure above shows slip 9954 which contains named user licenses, but slips can contain either named or concurrent user licenses.
- StarTeam named user licenses. This is the same as named user licenses in earlier StarTeam releases.
- StarTeam concurrent user licenses. This is the same as concurrent user licenses in earlier StarTeam releases.
- Unassigned. Select this "license type" when a user has no license.

Notes The status line at the bottom of the *User Manager* dialog provides licensing statistics.

You may need to edit the properties for the user to change the user name to the network logon name.

To assign a license to a new user:

1 From the *User Manager* dialog, select New User. The *New User Properties* dialog opens.

2 On the General tab, select a license type from the License drop-down list box. The default is StarTeam Concurrent.

3 Fill in the rest of the data on the General and other tabs as appropriate. Remember to use the network logon name for the User Name text box on the Logon tab.

4 Click OK to exit the *New User Properties* dialog.

Changing a User's Group Membership

A user can be a member of more than one group. If users belong to multiple groups, they can perform operations at the highest level permitted by any of their group privileges. For example, suppose that User A belongs to both the All Users group and the Administrators group and that the Delete Item privilege is granted to the Administrators group but *not* to the All Users group. User A can then delete any item in the server configuration projects.

Membership can be explicit or implicit. Membership in a group is explicit if:

- The group was selected at the time the user was created.
- The user's name was dragged to the group in a client
- The user's name was selected from the *Group Membership* dialog in the Server Administration utility.

Implicit membership is a result of the group hierarchy. If a user is a member of a child group, the user is also a member of the parent group, even if the member's name does not appear in the user list when you select the parent group. For a selected group that has child groups, you must select the Show Users in All Descendant Groups check box to see the complete list of members.

A user who is a member of a parent group and also a member of a child group within that group will have both implicit and explicit membership in the parent group.

This operation can be performed only when the server is running.

To review a user's explicit group memberships:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
   - Select Tools > Accounts > User Manager from the menu.

   These actions display the *User Manager* dialog.
4. Select the user. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:
   a. From the Groups tree, select the All Users group.
   b. Select the Show Users in All Descendant Groups check box.
5. Right-click, and select Properties from the context menu. The *User Properties* dialog appears.
6. Select the Membership tab. The list box displays the groups in which this user has explicit membership.

Tip To remove the user's explicit membership in a group, see "Removing a User from a Group" on page 89.

To change a user's explicit group memberships:

From the *User Manager* dialog:

1. Select the user from the User list.
2. Right-click and select Group Membership from the context menu. The *Group Membership* dialog appears.
3. Select the groups to which you want to add this user explicitly.
4. Click OK.

Removing a User from a Group

In most situations, removing a user from a group does not remove that user's account from the server configuration. However, if the user belongs to only one group, removing the user from the group is equivalent to deleting the user account.

This operation can be performed only when the server is running.

To remove a user from a group:

1. Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
   - Select Tools > Accounts > User Manager from the menu.

   These actions display the *User Manager* dialog.
4. Select the user. If the user's name does not appear in the Users list box, you can display a list of all users by doing the following:
   a. From the Groups tree, select the All Users group.
   b. Select the Show Users in All Descendant Groups check box.
5. Right-click the user's name, and select Remove from Group from the context menu. This action removes the user from that group only.
6. If the user belongs to only this group, the system displays the following message (*username* is replaced by the actual user's name):

   ```
   username is a member of only this group. Removing the user from the group
   would delete the user's account. Do you want to delete username user
   account?
   ```

7. Click Yes to delete this user account. Otherwise, click No.

Checking a User's Logon and Account Status

The Users list box in the *User Manager* dialog lists people by their user names, which are the names they use to log onto the application. The other columns of this list indicate:

- Whether the user is logged on and how many times the user has logged on.
- What the user's account status is.
- The user's full name.
- The date the user was added to the system
- The type of license assigned to the user

By using the scroll bar below the User list box, you can see all columns in the list.

Changes in a user's status can be performed only when the server is running.

To check a user's status:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 From the Groups tree, select the All Users group.

5 Select the Show Users in All Descendant Groups check box.

6 Review the information about the specific user that displays in the Users list box.

Tip To ensure that the information in the Users list box is current, click Refresh.

Removing a User's Account

Removing a user account removes that user from all groups in the server configuration. Although the user is removed, the user's name remains attached to any items assigned to or created by that user.

This operation can be performed only when the server is running.

To remove a user's account:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the user. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:

a From the Groups tree, select the All Users group.

b Select the Show Users in All Descendant Groups check box.

5 Right-click the user's name and select Delete Account from the context menu. The system displays the following message:

```
Do you want to delete username's user account?
```

6 Click Yes. This action permanently removes the user from the server configuration.

Suspending a User's Account

You may want to suspend user accounts for a number of reasons, such as keeping a workstation secure when someone is out of the office. To remove the suspension, you must reactivate the user account. See "Reactivating a User Account" on page 91 for more information.

This operation can be performed only when the server is running.

To suspend a user account:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the user. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:

a From the Groups tree, select the All Users group.

b Select the Show Users in All Descendant Groups check box.

5 Right-click the user's name and select Suspend Account from the context menu. The account status in the Users list box changes to "Suspended", and access to the server is denied after the user logs out.

Note You cannot suspend your own account.

Reactivating a User's Account

Sometimes a user account is been suspended by an administrator or inactivated because the password expired or a user is locked out due to failed logon attempts. In situations such as these, the user account must be reactivated for the user to have access to the application.

This operation can be performed only when the server is running.

To reactivate a user account:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the user. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:

a From the Groups tree, select the All Users group.

b Select the Show Users in All Descendant Groups check box.

5 Right-click the user's name and select Reactivate Account from the context menu.

These actions reactivate the user account.

Forcing a User to Log Off

You may have to force a user to log off for any number of reasons, including code violations or disaster recovery. When you force a user to log off, the user's next operation displays the following error message:

```
You are no longer logged on.
```

Depending on the reason for your action, you may need an additional method, such as e-mail or the telephone, to notify users to stop accessing the application.

To log on again, the user must exit the application and restart the client.

Most integrations between StarTeam and another application require the user to restart the application being used. However, these users are not usually notified that their connections to the server have been terminated.

This operation can be performed only when the server is running.

To force a user to log off:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the user. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:

a From the Groups tree, select the All Users group.

b Select the Show Users in All Descendant Groups check box.

5 Right-click the user's name and select Force Logoff from the context menu.The user is immediately denied access to the server configuration and to all projects residing in this server configuration.

Note You cannot force your own logoff.

Forcing a Password Change

It may be necessary to force users to change their StarTeam passwords if a project security breach has occurred. This operation can be performed only when the server is running.

To force users to change their passwords:

1 Do one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

These actions display the *Server Administration* dialog.

2 Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 From the *Server Administration* dialog, do one of the following:

- Click the Accounts bar in the lower left pane. Then click the User Manager shortcut.
- Select Tools > Accounts > User Manager from the menu.

These actions display the *User Manager* dialog.

4 Select the user. If the user you want to work with does not appear in the Users list box, you can display a list of all users by doing the following:

a From the Groups tree, select the All Users group.

b Select the Show Users in All Descendant Groups check box.

5 Right-click the user's name, and select Force Password Change from the context menu.

The Account Status column in the Users list box changes to "Password change required."

The user will be asked to change his or her password at the next log on. If the change is not made, the user is allowed access to the server configuration and the projects it contains, but will be locked out of the server configuration at the next log on. An error message warns the user that this will happen.

Note The accounts of users who fail to change their passwords can be reactivated by administrators. See "Reactivating a User's Account" on page 67

For the application, you can set the password expiration time limit, the minimum length, and require the use of strong passwords. These password properties apply to all user accounts on the server configuration.

Changes made to the password length properties take effect immediately, but apply only to new user accounts or new passwords. For example, if you change the minimum password length from eight characters to ten, all new users must have a password that is a minimum of ten characters long. However, existing users will still be able to use their eight character passwords.

Changes made to the expiration time limit take effect after the appropriate time interval. For example, if you change the password expiration time limit to thirty days, user accounts become are suspended if their passwords have not been changed before the time expires. Users will be prompted to change their passwords two weeks before the suspension takes place. The only user account not subject to expiration is the Administrator account.

By default, the strong password option is turned off. When this feature is turned on, as users change their passwords, they must provide strong passwords. Until such a change is made, their old “weak” passwords continue to work.

This operation can be performed only when the server is running.

To set password constraints:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the System Policy shortcut.
   - Select Tools > Accounts > System Policy from the menu.

   These actions display the *System Policy* dialog.
4. On the Passwords tab, select one of the following password expiration options:
   - Passwords never expire
   - Passwords expire after ___ days

     Enter the number of days a password will be valid.
5. Select the Require Strong Passwords check box to require passwords to meet all of the following criteria:
   - New password must be different from the old password.
   - New password must be different from the user name.
   - New password must be mixed case, containing at least one lowercase and at least one uppercase alphabetical character. (This is the English alphabet as determined by the ASCII value of the character.)
   - New password must contain at least one non-alphabetical character.

   Selecting this check box also changes the value in the “Minimum password length” text box to 3. You can increase it if you choose.
6. (Optional) Enter a minimum password length. The default, zero, allows passwords to be blank. The maximum password length is 32 characters.
7. Click OK.

Setting the Number of Logon Attempts

You can increase the security of your projects by entering a logon failure setting and duration. One cause of logon failure is hackers trying to figure out passwords for users. In such cases, you should consider changing the IP address of the system to make it more difficult for attackers to locate the server configuration and repeat their efforts. You may also want to change the user names of all users in the system.

You can configure the server configuration to notify members of the Security Administrators group by e-mail about logon failures and lockouts.

This operation can be performed only when the server is running.

Note See the notes for "Enabling E-mail Notification" on page 29 for information about changing the language of the messages sent by the server to the Security Administrators group.

To limit logon attempts:

1. Do one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
   - From a client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).

   These actions display the *Server Administration* dialog.
2. Select a server configuration from the list of servers. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. From the *Server Administration* dialog, do one of the following:
   - Click the Accounts bar in the lower left pane. Then click the System Policy shortcut.
   - Select Tools > Accounts > System Policy from the menu.

   These actions display the *System Policy* dialog.
4. Select the Logon Failures tab.
5. Select one of the following Logon Failure options:
   - Ignore

     This selection disables the logon failures option.
   - Lockout account after ___ failures

     Enter the number of logon failures you want to allow.
6. Select one of the following Lockout Duration options:
   - Forever

     With this option selected, only an administrator can reinstate the user.
   - Keep locked for ___ minutes

     Enter the number of minutes for the duration of the lockout. The user will be able to log on again after the designated timeout period.
7. To notify members of the Security Administrators group that users were locked out or attempted to log on unsuccessfully, select the By E-mail check box.
8. Click OK.

Reactivating Administrative Accounts

It is possible for any user, even users with an administrative account, to be locked out of a server configuration when the number of retries with the wrong password has been exceeded. The lockout period for the main administrative account (Administrator) is 24 hours. However, you can unlock the administrative account before the 24 hours have elapsed by using the following procedure:

To unlock the administrative account:

1. Shut down the server configuration, and disconnect its network connection to keep remote users off the server configuration.
2. Start the server configuration using the command line in foreground mode from the Server Installation folder.

   Example: starteamserver -start StarDraw -fg

   The configuration name specification is case-sensitive with the command line. The command prompt must be left open until the server configuration is shut down.
3. Set the system clock one day ahead.
4. Log in as Administrator and log off. This action will reactivate the Administrator account.
5. Set the clock back one day to its original time.
6. Shut down the server configuration by entering “X” and clicking Enter, which is how the server is shutdown in foreground mode.

Chapter

6

Managing Projects

A project is a high-level container that enables you to group related items (such as files) hierarchically. Views and folders enable you to organize those related items efficiently. For example, if you create a project for a software product, the files for the product's marketing requirements document, functional specification, source code, test suites, and user documentation can all be stored in separate folders.

Project views can be used so that developers see only the project's source code folder and its child folders, marketing personnel see only the project's marketing folder and its child folders, and so on. Each view can have a different folder as its root. Views also support branching and parallel development. At the view level or item-by-item, you can branch the files and other data. Branching enables your organization to create special variations of its products. For example, you can start on the 2.0 version of a product without interfering with the creation of service packs for the 1.0 version.

This chapter contains only the information about projects that administrators need to know. Refer to the *StarTeam User's Guide* for information on how to open an existing project, create desktop shortcuts to project views, or set data transfer options.

Before Creating a Project

Before you can create a project, a server configuration must have been created and be running. Each server configuration has one hive, which can store multiple projects. Although projects do not have to be related, items may be shared among projects using the same server configuration.

Most projects are based on a folder hierarchy that already exists on a workstation (i.e., hierarchy of working folders). For example, you may have created a set of files that compile into a product. The folder hierarchy represents the relationships between these files from a project point of view. It does not have to match the working folder hierarchy, but is often similar. Differences between the folders can arise if you:

- Do not add all the folders in the working folder to the hierarchy.
- Add folders to the hierarchy that are not in the working folder hierarchy (e.g., folders from another folder on your computer or from another project)

The folders on your computer that are used to create the hierarchy serve as the default working folders for the corresponding folders. When you add or check in files, the application copies the files from the working folders to the repository. When you check

files out, the application copies the desired revisions of the files from the repository to the working folders. The files in the working folders are your working files.

After Creating a Project

After you create a project, the left pane of project view window displays the hierarchy of folders in that project. Initially, all files in the project folders will have a status of “Not in View.” You will need to add the appropriate files to the project and create any additional views you want for this project. For more information about files and views, see “Creating Additional Views” on page 144 in this guide and related material in the *StarTeam User’s Guide*.

In the view window shown below, StarDraw is the root folder. (It has the same name as the project and the project view—both of which appear in the window title bar.) User Manual is the selected child folder. Because the User Manual folder and the File tab are selected, the files associated with the User Manual folder are listed in the upper pane.

Because Reference Manual.doc is selected in the upper pane and the Detail tab is selected in the lower pane, the lower pane displays details about that file, including its name, status, revision, time stamp, and size.

Adding Access to a Server Configuration

Before you can create a project, the server configuration in which the project will be placed must be running and you must be able to access that configuration from your client workstation. To add access to a server configuration, you must provide its domain name (DNS) or IP address. After the server configuration has been added, you can create a project on it or open existing projects that are available on the configuration.

To add access to a server configuration:

1. On the Server, verify that the configuration is running.
2. Launch a client, and do one of the following:
   - Click the New Project icon on the toolbar. The *New Project Wizard* dialog appears.
   - Click the Open Project icon on the toolbar. The *Open Project Wizard* dialog appears.
   - Select Tools > Server Administration from the menu bar. The *Server Administration* dialog appears.
3. Click Add Server. The *StarTeam Server* dialog appears.
4. Enter a description in the Server Description text box. This description serves as a unique name for the server. It is *not* case-sensitive and cannot contain colons (:).
5. Enter or browse for the address in the Server Address text box. The address is the computer name or IP address.
6. Enter the endpoint in the TCP/IP Endpoint text box. The endpoint is the port number associated with the protocol.
7. (Optional) Select the Compress Transferred Data check box to compress data transferred between your workstation and the server.
8. (Optional) Select a minimum encryption type check box, to encrypt data transferred between your workstation and the server.

Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines. For more information about encrypting and compressing transferred data, see the *StarTeam User's Guide*.

The encryption types are ordered (top to bottom) based on speed and security. Each encryption type is slower, but safer, than the type that precedes it in the list.

9 Click OK.

Creating a Project

You can create a project on any server configuration for which you have access and the access rights required to a create project. If your client is not currently set up to access the server configuration on which your project will reside, you must add access to that server configuration as part of creating the project.

When you create a project, you must provide a project name and designate a working folder location for the project root folder. For more information about folders, see the "Managing Folders" chapter in the *StarTeam User's Guide*.

The initial view of a project is created when you create the project. It has the same name as the project, although you can change it later if you choose.

Note To create a project with existing Visual SourceSafe (VSS) or PVCS files, see Chapter 13, "Using StarTeam with Visual SourceSafe," on page 219 or Chapter 14, "Using StarTeam with PVCS," on page 229.

To create a project:

1 Open a client and do one of the following:

- Click the New Project icon on the toolbar.
- Select Project > New from the menu bar.

The *New Project Wizard* dialog appears.

2 Do one of the following:

- Select the server on which you want to store your project, and click Next to continue.
- Click Add Server to add access to a server that does not appear on the list, and follow the steps in "Adding Access to a Server Configuration" on page 74. Then select the server configuration.

3 If you have not yet logged onto the server configuration, you will be asked to do so. After you log on, the *New Project Wizard* dialog appears.

If you are already logged onto the server configuration or are running the Toolbar, you have the option of logging on as a different user. To do this from the *New Project Wizard* dialog, click the Log On As button, and enter the alternate user name and password in the *Log On* dialog. To reset your default credentials, select the Save As Default Credentials for This Server check box. When you are done, click OK.

4 On the *New Project Wizard* dialog, click Next. This action displays the *New Project: Project Name* dialog

a Enter the name of the project in the Project Name text box.

b Enter a description of the project

c Click Next. The *New Project Wizard: Working Folder* dialog appears.

5 On the *New Project Wizard: Working Folder* dialog:

a. For the project root folder, enter or browse for the location of the default working folder. If the project root folder does not yet exist, the application will create it for you.

b. Click Next. The *New Project Wizard: Child Folders* dialog appears.

6. (Optional) On the *New Project Wizard: Child Folders* dialog, exclude or include child folders in the working folder from the folder hierarchy on a folder-by-folder basis.
   - Select the folders to be removed from the project, and click Exclude.
   - Click Reset to re-list excluded folders.

   You can include additional folders after a project is created.
7. Click Finish to open the project.

Creating a Project by Dragging a Folder

On Windows systems, you can create a project by dragging folders from Windows Explorer to a client.

To create a project using drag-and-drop:

1. From My Computer or Windows Explorer:
   a. Using the mouse, select the folders to be included in the project.
   b. Drag the folders to the left pane of an open view window.
   c. Release the mouse button. the application adds the folders to that view of the project.

   The *New Project Wizard* dialog appears.
2. Use the *New Project Wizard* to create the new project. See “Creating a Project” on page 75 for instructions.

Opening an Existing Project

Before you can open a project, you must have access to the server configuration on which the project resides. See “Adding Access to a Server Configuration” on page 74 for more information.

To open a project window:

1. From a client, do one of the following:
   - Click the Open Project icon on the toolbar.
   - Select Project > Open from the menu bar.

   The *Open Project Wizard* dialog appears.
2. Click the plus sign in front of the server name, or double-click the name of the server on which the project is located.
3. If you have not yet logged onto the server configuration, you will be asked to do so. After you log on, the *Open Project Wizard* dialog displays.
4. If you are already logged onto the server configuration or are running the Toolbar, you have the option of logging on as a different user. To do this from the *Open Project Wizard* dialog, click the Log On As button, and enter the alternate user name and password in the *Log On* dialog. To reset your default credentials, select the Save as Default Credentials for This Server check box. When you are done, click OK..
5. From the *Open Project Wizard* dialog, do one of the following:

- Open the root view of the project by selecting the project and clicking Finish.
- Select a specific project view by double-clicking the project name or by selecting the project name and clicking Next.

  Either of these actions displays the *Open Project Wizard: Select View* dialog. Select a view name from the View list, and click Finish to open the project in that view.

Note If you see a red X on a View icon, you do not have access to that view.

Modifying Project Names or Descriptions

You can use the *Project Properties* dialog to review or change the project name and description, enable keyword expansion, or force users to provide a reason for checking in a file.

To change the project name or description:

1. From an open project in a client, select Project > Properties from the menu bar. The *Project Properties* dialog appears.
2. Select the Name tab.
3. Enter a new project name in the Name text box.
4. Enter a new project description, or edit the description in the Description text box.
5. Click OK.

Using Keyword Expansion

If keyword expansion is enabled for a project, you can embed keywords within a project text file. The keywords are automatically expanded during file check-out to provide file and revision information within the file. For a list of the keywords available in the application, see the *StarTeam User's Guide*.

When you put keywords in file, use only one keyword per line. Keywords can be used in plain text files, but are not supported in files that use an .rtf text format. Keywords are not expanded for Unicode files.

To enable keyword expansion:

1. Select Project > Properties from the menu bar. The *Project Properties* dialog appears.
2. Select the Options tab.
3. Select the Keyword Expansion check box to enable keyword expansion and use keywords in your text files.

   This check box applies to files added or checked in from a client, or from integrations with third-party applications.
4. Enter the file specifications with which you want to use keywords. If this text box remains blank, keywords are not expanded.

   You can use wildcards within the file specifications. For example, entering *.bat, *.cpp (without quotation marks) indicates that keyword expansion should be used with all files having either of those two extensions. The file specifications list can contain a maximum of 254 characters.

   You can use a space, comma, or semicolon as the keyword delimiters. To include a delimiter (such as a space) in a file specification, enclose the file specification in double quotation marks.

5 Click OK.

Forcing Users to Provide a Revision Comment

Normally, if you use the File drop-down menu or context menu to check in files, the *Check In* dialog appears. This dialog provides a text box in which you can supply a reason for checking in the files.

Normally, if you click Check In (black arrow) or Check In and Unlock (blue arrow) from the toolbar, you perform what is called a “fast” check-in, in which the *Check In* dialog does not appear.

However, if you force users to always supply a reason for creating a new revision, the *Check In* dialog always opens and the application requires you to fill the Comment text box, regardless of whether you use the toolbar, menu, or context menu.

To force all users to supply a check-in reason:

1 Select Project > Properties from the menu bar. The *Project Properties* dialog appears.

2 Select the Options tab.

3 Select the Require Revision Comment When Files Are Checked In check box.

4 Click OK.

Note You cannot force users to supply a reason for checking in files when they are using an integration with a third-party application.

Forcing Users to Lock Files Before Check-In

By default, the application does not force users to lock the files that they plan to modify. However, depending the policies and processes of your company, you may want users to exclusively lock files before modifying them. If that is the case, you can use the application to stop users from checking in files that have not been exclusively locked.

Requiring users to exclusively lock files that they plan to change helps to avoid situations in which two or more users make changes to a file and the changes must be combined. In such situations, the files must be compared and/or merged.

To help prevent the need for merges, users should also:

- Notice whether a file is already exclusively locked by another user before checking it out to modify it.
- Lock each file before they change it, to alert other users to their intentions.
- Verify that each working file has Current status, to avoid making changes to older revisions of the files. If a file does not have Current status, it must be checked out before any changes are made.

To force users to lock files before checking them in:

1 Select Project > Properties from the menu bar. The *Project Properties* dialog appears.

2 Select the Options tab.

3 Select the Require Exclusive Lock When Files Are Checked In check box. When this option is selected, only the user who exclusively locked a file can check it in.

4 Click OK.

Note This project option overrides the Use Non-exclusive Locks in Integrations option on the File tab of the *Personal Options* dialog.

Making Unlocked Files Read-Only

In many cases, users make edits before realizing that their files must be exclusively or non-exclusively locked before they can be checked in. If the files are read-only, users are less likely to make this mistake.

To make working files read-only when they are unlocked:

1. Select Project > Properties from the menu bar. The *Project Properties* dialog appears.
2. Select the Options tab.
3. Select the Mark Unlocked Working Files Read-only check box. With this selection, working files that change from locked to unlocked become read-only when users perform the following file operations:
   - Check-in
   - Check-out (from file or history list)
   - Unlock

   As a consequence, only locked files can be edited.
4. Click OK.

Notes This project property overrides the identical Mark Unlocked Working Files Read-Only personal option. When selected, it overrides the personal option. It applies to files that are unlocked either from a client or from integrations with third-party applications.

If you select this project property (or the personal option) and then change your mind, verify that no files are read-only *before* clearing the check box. Next, force a check out and lock all the files (or just the read-only files) Finally, unlock them.

If this check box is cleared, you must use the operating system to change the read-only attribute to read/write.

Understanding Process Rules

Process rules take change management to a higher level by requiring developers to follow a defined development process which ensures that all changes are linked to either a change request, requirement, or task. Items used in this way are known as process items. Thus, code and content are modified only as the result of a clearly defined and approved objective.

The application's process rules are especially useful when creating baseline builds or configurations. A build is a labeled configuration that identifies the file revisions and process items defining the code and content baseline. Without process rule enforcement, developers using the application can create baselines either by:

- Labeling an entire project view at a point in time.
- Associating file revisions with a revision label on check-in.

Because process rules require each new file revision to be linked to a process item, the development team can promote changes into baselines by:

1. Starting with the previous baseline (for example, by checking it out based on its label).
2. Selecting process items for inclusion in the new baseline.

   For example, custom build utilities available from Borland Professional Services (at www.borland.com/services/) can use the pinned links between file revisions and the selected process items to identify which file revisions to include in the new baseline.

As a result, development teams can better control the changes made to these baselines.

3 Labeling the new baseline.

Even without the enforcement of process rules, a user can voluntarily use a process item when adding or checking in files. In addition, the user can assign that process item a specific status (e.g., a change request can be marked as Fixed). For convenience, a user can also select a change request, requirement, or a task as the Active Process Item before to adding or checking in files. See the *StarTeam User's Guide* for more information.

Requiring the Use of Process Items

If you have the access rights necessary to change project properties, you can:

- Require the use of process items.
- Stipulate that only specific types of items with specific statuses can be used as process items.

In addition, to use process items, you must ensure that project users have the necessary access rights, which are the rights to:

- See and modify the types of items used as process items in the project view.
- Create and modify links on files and process items.

To enforce process rules:

1 Select Project > Properties from the menu bar. The *Project Properties* dialog appears.

2 Select the Process Rules tab.

3 Select the Require Selection of Process Items When Files Are Added or Checked In check box.

4 To allow users to use a change request as a process item:

   a Select the Permit Selection of Change Requests as Process Items check box.

   b Specify what change requests will be available:

      - Select the Open check box to use only change requests with the Open status as process items.
      - Select the In Progress check box to use only change requests with the In Progress status as process items.
      - Select both the Open and the In Progress check boxes to use only change requests with either of these statuses as process items.
      - Select neither the Open nor the In Progress check boxes to allow any change request to be used as a process item, regardless of its status.

5 To allow users to use a requirement as a process item:

   a Select the Permit Selection of Requirements as Process Items check box.

   b Specify what requirements will be available:

      - Select the Approved check box to allow only approved requirements to be used as process items.
      - Clear the Approved check box to allow any requirement to be used as a process item.

6 To allow users to use a task as a process item:

   a Select the Permit Selection of Tasks as Process Items check box.

b Specify what tasks will be available:

- Select the Ready To Start check box to use only tasks with the Ready To Start status as process items.
- Select the In Progress check box to use only tasks with the In Progress status as process items.
- Select both the Ready To Start and the In Progress check boxes to use only tasks with these statuses as process items.
- Select neither the Ready To Start nor the In Progress check boxes to allow any task to be used as a process item.

7 Click OK.

Using APEs with Projects

If you use Enterprise Advantage, you can use alternate property editors (APEs) for a project. APEs are Java forms created specifically for your company in support of a corporate process. Workflow processes are created for use with the forms. APEs can be created for the File, Change Request, Requirement, Task, and Topic components and can be defined for one or more components. APEs use StarTeam Runtime to access the Server. They can be customized because they are implemented in standard programming languages. Sample property editors for several application components are included. For more information about APEs, see the *StarTeam Extensions User's Guide*.

You must define an APE for any project that will use the custom forms and workflow configurations.

The Extensions have special executable files, Locator.exe and LocatorDebug.exe, to ensure that the latest workflow configurations and custom forms are used on the client workstation.

Locator.exe and LocatorDebug.exe are small executable files that use the bootstrap.jar file. All three files are found in the client's installation folder.

Important Be careful when setting the APE for a project that is in production because the new setting takes effect immediately. It is important to test any changes that are made to an APE before making this editor available to a wide audience.

To use an alternate property editor (APE) with an application component:

1 Select Project > Properties from the menu bar. The *Project Properties* dialog appears.

2 Select the Editors tab.

3 Select the appropriate Use Alternate Property Editor for <item> check box.

4 Enter or browse for the correct APE. Usually you enter the word Locator followed by the name of the APE.

Locator is a program on each client workstation that distributes code and XML updates to client workstations. If specified here, whenever the application requests the APE, Locator looks for the Extensions project managed by the server configuration. If that project exists and contains an APE with the specified name, Locator copies the tip revision of that APE and its related files to the client workstation (if they are *not* already there). Then the user uses the specified APE instead of the standard dialog for the specified type of item.

5 Click OK.

Deleting Projects

To delete a project, you must have this privilege or access right. Before you delete a project, be absolutely certain that you want to perform this action, because you will not be able to access any item in a deleted project that is not shared with another project. After a project is deleted, it is not be visible in the *Open Project* or the *Select View* dialogs. If other users are connected to the project/view when it is deleted, they will receive a “deleted” message the next time they initiate a project or view command.

Note Deleting a project does not remove its contents from the server configuration database. However, unshared items will no longer be accessible.

To delete projects:

1. Select the root project folder.
2. Select Project > Delete. A confirmation dialog asks you to confirm the deletion.
3. Click Yes to confirm. A second confirmation dialog asks you to enter the name of the project.
4. Enter the name (which is case-sensitive) in the Project Name text box.
5. Click OK. This action closes the open view window for the project. To display a view from another project, follow the instructions in “Opening an Existing Project” on page 76.

Assigning Project Access Rights

Administrators can assign project-level access rights. See Chapter 6, “Managing Projects,” on page 73 for details.

Chapter

7

Using Promotion States

Most projects go through a release or production cycle. For example, a software application often cycles from developers to testers to the marketplace. Then work starts on the next release.

Promotion states provide a convenient mechanism for ensuring that the right files (or other items) are available to the right people at the right stage of the production cycle. For example, if a software administrator creates the promotion states Test and Release, the files that are being worked on by testers can be assigned to the Test state and the files that are ready for release can be assigned to the Release state.

The promotion state feature permits an administrator to create promotion states and associate a view label with each state. The view label for a state can be changed whenever appropriate. It can also be promoted from one state to the succeeding state. For example, although testers may always be using files in the Test promotion state, the files may be from Build 07 in one week and from Build 08 in the next.

Notes Many features of the application depend on calculations involving times and dates. In particular, labels, configurations, and promotion states are all governed by time and date calculations. If the clients and the Server are not kept synchronized, a number of operations (such as checkouts, file status displays, or label creation) may fail or produce inaccurate or unreliable results.

How an Organization Might Use Promotion States

Suppose that a software company wants to use the following promotion states to correspond with its use of the application:

- Development

  Developers work with the tip revisions of files. These files have no view label because they are constantly changing. Many companies do not use Development as a promotion state, because configuring a view to a promotion state, even when the view label for the state is <current>), makes the view read-only.

- White Box Test

  Testers check both the source code and the compiled executables for problems that need to be fixed. The source code will have a view label to ensure that the testers are looking at an unchanging set of files. The view label will be attached to a

promotion state named White Box Test. (White box testing is testing with full knowledge of what is in the source code.)

The executables are not stored in the application because they can be easily built from the source code. Testers install them from a Builds folder on the network. This folder has child folders named Build 1, Build 2, and so on.

Change requests are entered against both the source code and the executables. Developers make repairs in the current source code, sometimes reviewing the files with the view label attached to the White Box Test promotion state.

- Black Box Test

  Testers install the executables, just as they would with white box testing. However, they do not need to see the source code or use promotion states with it. (Black box testing is testing with no knowledge of what is in the source code.)

  Change requests are entered against the executables only. Developers make repairs in the current source code, sometimes reviewing the files with the view label attached to the Black Box Test promotion state.

- Alpha

  End users of the software product being developed install the product executables and test the product in their own environments.

  Change requests are entered by the alpha coordinator and/or the users against the executables only. Developers make repairs in the current source code, sometimes reviewing the files with the view label attached to the Alpha promotion state.

- Beta

  Beta testing is similar to alpha testing, but the group of users is greatly expanded because the product is much more stable.

  Change requests are entered by the beta coordinator and/or the users against the executables only. Developers make repairs in the current source code, sometimes reviewing the files with the view label attached to the Beta promotion state.

- Release

  The product is now sold in the marketplace. Users install the executables and call product support. Product support enters change requests against the executables only. Developers make repairs in the current source code, sometimes reviewing the files with the view label attached to the Release promotion state.

  The fixes go into future product releases and service packs to releases already in the marketplace.

In this example, every time the source files are used to produce a build (set of executables) for testing, a view label is applied to the files to identify them for future reference. It is convenient to use view labels such as Build 1, Build 2, etc. so that it is clear which source code files were used to create which set of executables.

Over time, the build or view label associated with a promotion state will change. For example, the Release state may initially be associated with <current>, rather than a view label, because no files are candidates for release and no appropriate view label has been created.When white box testers decide that the set of files that they have examined is ready for black box testing, the view label associated with the White Box Test promotion state will be moved to the Black Box Test promotion state, and so on.

If promotion states are used, developers and testers who look at source code do not need to know that view label Build 120 is currently being checked by white box testers, that the executables for Build 117 are currently undergoing black box testing, and other details.

Creating Promotion States

To create promotion states, you must have the access rights that allow you to do so. For more information on this subject, see “Understanding Access Rights for Promotion States” on page 212.

When creating promotion states, many administrators assign <current> (instead of a view label) to the initial promotion state because tip revisions will be always be used by that state. They may also select <current> for later promotion states that currently have no view labels. They assign a view label to a state when the files associated with that label meet the criteria for that particular state. If desired, a view label can be promoted from a preceding promotion state to a label-less state.

To create promotion states:

1. Select View > Promotion from the client menu bar. The *Promotion* dialog displays any states currently created for the view.

   Promotion states display from last to first. For example, if a file goes through White Box Test, followed by Black Box Test, Alpha, Beta, and finally Release, you must list these states in the *Promotion* dialog in reverse order: Release, Beta, Alpha, Black Box Test, and then White Box Test.

2. Click Add. The *Promotion State* dialog appears.
3. Enter the name of a promotion state in the Name text box. The name you enter should represent the initial promotion state attached to the set of files.
4. Enter a description for the promotion state in the Description text box.
5. To assign a view label to the state, select it from the View Label drop-down list box.

   Existing view labels are listed in reverse chronological order based on the time for which they were created. You can change this label later by using this dialog or by promoting it to the next state.

6. Click OK.
7. Repeat steps 2 through 6 for the following promotion states. The states should appear in ascending order from the final state to the initial state.
8. When you are done, click Apply on the *Promotion State* dialog. (Clicking OK on the *Promotion State* dialog does not save your changes.)

Tips

- If you enter the promotion states in the wrong order, use the Move Up and Move Down buttons to rearrange them.
- Click Reset to restore the list of promotion states available when you opened the dialog or last clicked Apply.
- Click Edit to change the name, description, or view label of a selected promotion state.
- Click Delete to delete a selected promotion state from the list.

Promoting a View Label

You can promote a view label from one promotion state to the next if you have the appropriate access rights.

Suppose that your promotion states are White Box Test, Black Box Test, Alpha, Beta, and Release and that Build 17 is the view label assigned to Black Box Test.

When Build 17 has the correct approvals, you can promote it from Black Box Test to Alpha. The promotion process assigns Build 17 to Alpha, replacing the view label previously assigned to Alpha.

However, at least temporarily, Build 17 is still the view label assigned to the Black Box Test promotion state. If you want to alter this, you can manually edit edit the promotion state or promote the Build 18 view label from the White Box Test promotion state to Black Box Test.

To promote a view label from one state to the next:

1. Select View > Promotion from the client's menu bar. The *Promotion* dialog displays the promotion states currently created for this view. The states display from the final state down to the initial state.
2. Select the promotion state currently associated with the view label that you want to promote.
3. Click Promote. The *Promote View Label* dialog indicates that the view label is now associated with the next state
4. Verify that the information is what you were expecting to see, then click OK. The view label applies, at least temporarily, to both the state it was promoted from (in this example, Black Box Test) and the state it was promoted to (Alpha). Usually, your next action will be to associate a new view label with the Black Test promotion state.
5. When you are done, click OK.

Demoting a View Label to a Previous State

Sometimes a labeled set of files is promoted prematurely and must be demoted. For example, if Build 17 is promoted to Alpha, but outside users are encountering serious problems with it, then Build 17 should be discarded or returned to the Black Box Test promotion state.

View labels can be demoted only by editing the promotion state. For example, to demote the Build 17 view label, you would select Alpha from the *Promotion* dialog and associate it with the view label that it had prior to Build 17. If necessary, you would then select the Black Box Test promotion state and the Build 17 view label to it.

To demote a view label:

1. Select View > Promotion from the client menu bar.
2. Select the promotion state to be edited.
3. Click Edit. The *Promotion State* dialog appears.
4. Change the view label by making a different selection from the View Label drop-down list box.
5. Repeat steps 2-4 for other promotion states, if required.
6. Click OK.

Basing a Configuration on a State

When you configure a view, a folder, or an item, you can select a promotion state.

Whether the promotion state is assigned to <current> or to a specific view label, a view configured to a promotion state is read-only. This status can create problems for some users, such as white box testers.

If testers configure a view (roll it back) to the White Box Test promotion state , they can access and test the correct source code files, but they cannot add change requests to the view, because the roll-back configuration is read-only. (Change requests can be added to the current configuration only.)

In this situation, testers may need to open two instances of a client, one with the view configured to the White Box Test promotion state and the other using the current configuration.

Black box testers or alpha and beta coordinators will not encounter problems when a view is configured to a promotion state, because they do not need to look at source code files. They test installed executables, so they can enter change requests into the view current (Development) configuration.

To base a view configuration on a promotion state:

1. Select View > Select Configuration from the client's menu bar. The *Select A View Configuration* dialog appears.
2. Select the Promotion State Configuration option button.
3. Select a state from the drop-down list box.

   Configuring a view to a promotion state maintains the dynamic nature of the promotion state. If the promotion state's view label changes, the view is configured to the new view label on the next refresh. The promotion state's name appears on the status bar. If you exit the view and return, you return to the current configuration.
4. Click OK.

To base a folder or item configuration on a promotion state:

1. Select the Behavior command from the appropriate menu. The item *Behavior* dialog appears.
2. Select the Configuration tab.
3. Select the Promotion State Configuration option button.
4. Select a state from the drop-down list box.

   Configuring a folder or item to a promotion state is *not* dynamic. When you select a configuration for a folder or item, the application sets (or pins) the folder or item to the revision with the promotion state's current view label. This configuration does not change when the promotion state changes. If you return to the dialog in which you selected the configuration, you can see the date and time of that revision as the configuration setting.
5. Click OK.

Note A folder that has been configured to a past state can be identified only by looking at its Configuration tab. See the *StarTeam User's Guide* for information on locating items that have been configured to a past state.

Setting Access Rights for Promotions

Access rights can be set for a project or a view that allow users to:

- Create, delete, and reorder promotion states.
- Edit the name or label for a state.

Access rights can also be set for individual promotion states that allow users to:

- Edit the view label for the state.
- Promote the state.
- Set access rights for the state.

For more information, see "Understanding Access Rights for Promotion States" on page 212.

Chapter

8

Using View Manager

On the Cross-Platform client, you can use View Manager to compare and merge views. View Manager provides the functionality needed to copy and synchronize files and related data, such as labels, between views.

View Manager can be run as an application with a graphical user interface or as a command-line utility. The command-line utility is faster, but the graphical user interface provides a comparison tree that enables you to preview what will be done or verify what has been done.

The functionality of the two approaches varies slightly. For example, if you are using the graphical user interface, manual file merging is postponed until the end of the run operation. From the command line, you can either postpone manual file merging or display a merge utility whenever View Manager finds a pair of files to be merged.

With View Manager, you can:

- Copy a folder from one view to another.
- Copy a folder as a child of another.
- Synchronize two folders from different views.
- Copy views from one project to another. The source view can be configured to a particular label or promotion state.
- Copy a view as the child another.
- Synchronize two views.

Operations on folders can include their child folders., and operations on views can include their child views. However, the larger the amount of data, the longer the process takes.

To use View Manager, you must specify a source object, a target object, a set of options, and an operation, such as copy-to-target or source-to-target. View Manager then copies file histories, labels, promotion states, and information about who created these objects, depending on the options selected. Most properties, such as descriptions, are replicated along with the objects.

Preparing to Use View Manager

Before using View Manager, you must prepare the operation. The steps you should take apply regardless of whether you intend to use the View Manager Graphical Interface or the command line.

Minimize What Must be Replicated

Pin-point the data that must be replicated as exactly as possible. If you do *not* really need some objects, do *not* include them.

Understand View Manager Capabilities

Before using View Manger, review how the application works so that you are become familiar with the issues that may arise during a merge. In most situations, workarounds exist. Regardless of the type of merge you intend to perform, Borland strongly encourages you to backup the server.

Replication of Files

The only items that View Manager replicates files. No change requests, requirements, topics, or tasks are copied or synchronized.

Replication of Reference Views

View Manager cannot copy a reference view without also copying its parent, because they access the same files. For example if a reference view is a child of a root view, you must copy both the root view and the reference view. If the parent is another reference view, you must use the grandparent or an even earlier ancestor.

View Manager can select an appropriate parent for you. When it does, the Include Child Views option applies to the initially selected view (the reference view) rather than to the parent view.

Replication of Foreign Archives

Foreign archives are not supported by View Manager. If projects were created by wrapping PVCS and VSS files, that project's files must be converted to native format before being replicated. If a PVCS or VSS file is shared into an otherwise native project, that file is bypassed and an error is added to the log file.

Creation Time Issues

Although View Manager can give files, folders, views, and projects dates in the past, it actually creates those objects in the present.

Thus, if View Manager adds objects that are older than their new parent objects, it gives the objects the current time because the server will not accept the original time. For example:

- Views cannot be older than their parent views or the project.
- Child views cannot be older than their parent views.
- Root folders cannot be older than their views.
- View labels cannot be older than the view itself.
- Promotion states cannot be based on view labels that are older than the view.

Most problems with object creation times can be avoided by performing the initial copy-to-target operation or source-to-target synchronization with an empty target project, view, or folder. When the target object is empty, View Manager deletes it and attempts to create a new one with the same name and an appropriate creation time in the past.

- A project is considered empty if it has only a root view and that root view has only a root folder and no files.
- A view is considered empty if it is either a root view or a blank view and contains only a root folder and no files.
- A folder is considered empty if it is the root folder in either a root view or a blank view and contains no child folders or files.

Re-creation always works for empty projects, empty root views, and empty root folders in root views. However, if the empty target object is a blank view or the root folder of a blank view, its time cannot be made earlier than the creation time of its parent view.

Replication of Floating Branching Views

If a child view is configured to branch and float, files and folders added to the parent view after child view has been created are also propagated to the child view. After they are in the child view, they may branch when they change.

However, View Manager cannot exactly replicate this behavior.

If the newer file has *not* branched in the child view, View Manager shares it from the parent view into the child view.

If a newer file has branched in the child view, View Manager creates a new object for the file.

For a newer folder, View Manager creates a new folder, regardless of whether it has have branched or not.

Thus, the references are not accurate and the new files have no file history. Files in the newer folders are not treated any differently from newer files in other folders. They are shared into the child view unless they have branched.

Replication of Non-Floating Branching Views

If you are replicating a non-floating branching view (that is, a branching view whose initial configuration was set to a specific time, view label, or promotion state), you *must* replicate the creation times to keep the view a non-floating view. Otherwise, View Manager creates the replicated view as a floating branching view (that is, a view with an initial configuration of current).

Replication of View Labels and Promotion States

View labels and promotion states that have the same names in both the source and target view are never copied. Therefore, if Build 1 is a view label in both locations, that label can never be changed in either location by View Manager.

This feature can be good or bad. If the two views coincidentally have different labels with identical names, you would not want either of them overwritten. However, if Build 1 is in one view solely because it was copied from the other view, you might want View Manager to update information about that label.

To avoid problems, Borland recommends that view labels be maintained in only one of the two locations. Be aware that you can use the application to change the name of a view label or delete it if necessary.

Replication of Child Views

If the source or target object is the parent of at least one child view that will also be copied, you must select the Copy File History option. (This option is selected by default.)

If you copy a source view to a target view that is unrelated to it, the view is copied as a blank view and all the files are recreated. For example, if you copy a root view as the child of another project's root view, all the files are considered new files. If you copy a branching view from one project as a child of a view in another project, all the files are considered new files.

Replication of Tip Revisions

If you copy only tip revisions, the tip revision of each file becomes revision 1 in the target view. The revision numbers will not be the same in the source and target objects, but that will not cause any synchronization problems.

However, if you create a child view at a later time, you may prefer to have the file history. When you replicate a view along with at least one of its children, View Manager requires you to copy the file history.

File Name Changes

Avoid changing object names if you use View Manager. If you have changed the name of a file in a folder before you replicate that folder, the file will show up as two files in the target location—one with the old name and one with the new name.

Unfinished Operations

If you stop View Manager or it aborts a run operation, the work that has already been done is *not* undone. In this situation, you can do one of the following:

- Resume the operation.
- Use the application to delete the partially completed projects, views, and so on.

Using the View Manager Main Window

The main window for the View Manager graphical interface has three panes. The two upper panes, called the left and right panes, display the project-view-folder tree for a specific server configuration. You will select the source and target objects from these

trees. The lower pane contains a combined tree and list control that is explained on page 96.

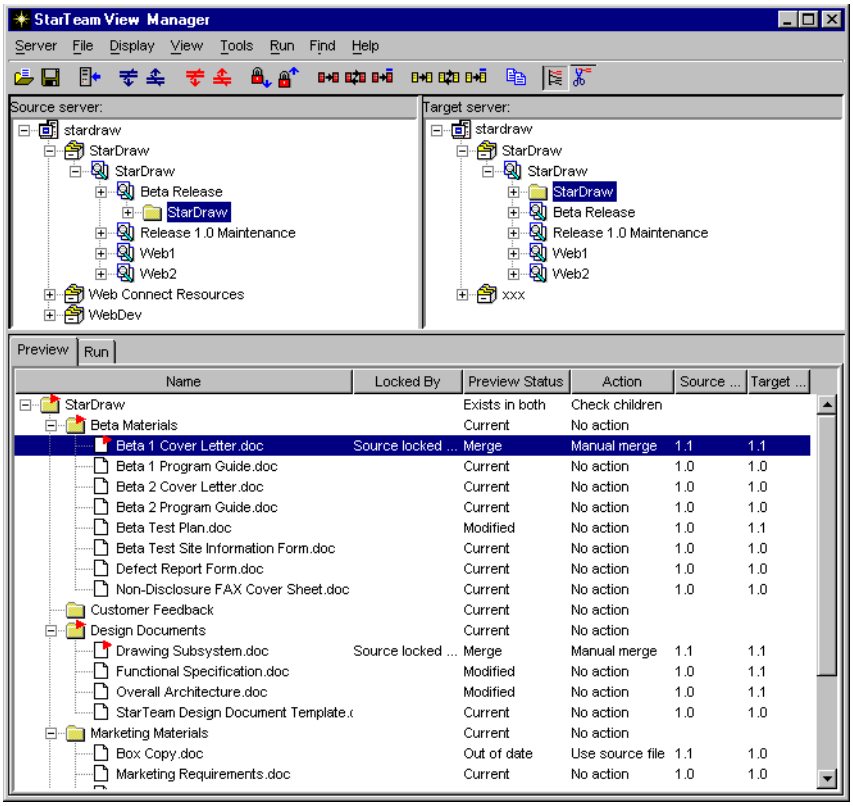

The server, projects, views, and folders are nodes on the trees. Icons indicate the type of node.

When you specify the use of a merge utility to resolve the differences and/or conflicts in text files, the corresponding files are represented by the flagged file icon. Use of the merge utility is postponed until the portions of the run that can be performed automatically have completed. Unresolved binary files are also flagged, but cannot be resolved by using the merge utility for text files.

By design, projects have no child projects. Each project has only one immediate child (its root view). Each view can have several immediate children that are views, but only one immediate child that is a folder (its root folder). The root folder is displayed as the first of the view's immediate children.

You can expand or collapse branches of a tree:

- A plus sign identifies a collapsed branch. Clicking it expands the branch.
- A minus sign identifies an expanded branch. Clicking it collapses the branch.
- Double-clicking the names of the server, projects, views, or folders in the tree expands or collapses branches.

Using the Upper Panes

The upper panes show the tree for the selected server configuration. You can select or change the source and target objects by using menu commands or open a properties file, which contains a pair of objects and options saved in the past.

First, you select a view or folder from the left pane. The right pane must contain the same type of object as the left pane or, in the case of a folder, an object *one* type

higher. For example, if a folder is selected from the left pane, you can select either another folder or a view from the right pane.

When you select a view, the merge process applies to that view's folders or to its folders, child views, and their folders, depending on the selected options..

When you select a folder, the merge process applies to that folder's files only or to its files, child folders, and their files, depending on the selected options.

After choosing objects from both source and target panes, you can run the following operations:

- Preview. Displays the results as they would occur in a source-to-target, bidirectional, or missing-objects-only operation.
- Source-to-target synchronization. Corresponding files can be merged and checked into the target. Missing objects can be added and deleted. Specific details depend upon the option settings.
- Bidirectional synchronization. Corresponding files are merged and can be checked in to both the source and the target. Missing objects can be added and deleted. Specific details depend upon the option settings.
- Missing-objects-only synchronization. Missing objects can be added and/or deleted, depending upon the option settings. No merging takes place.
- Copy to target. The source object is copied as a child of the target object. No merging takes place. If you have selected a folder from the left pane and a view from the right pane, the copy-to-target operation is the only run operation available.

If the source and target objects are the same type but have different names, they are treated as though they were equivalent during a synchronization operation. For example, if you select Folder X as the source and Folder Y as the target and Source To Target at the operation, View Manager treats Folder X and Folder Y as though they were both named Folder Y. When views are equivalent, their root folders are considered equivalent, regardless of their names.

The following tables explain source and target object selections and the results of the operations performed on them. To compare two projects, you must select their root views and use the Include Child Views option.

Table 8.1 Source Folder to Target Folder

Operation	Results
Preview	The lower pane shows the results of performing a source-to-target, bidirectional, or missing-objects-only synchronization, given the currently selected folders and options. The preview starts with the two corresponding folders and ignores the rest of the tree. Child folders may also be included, depending on the selected options.
Source-to-target Sync	All the files in both folders are synchronized. Merged files can be checked into the target. Child folders that exist in both folders may also be synchronized, depending on the selected options. What happens to objects that exist only in the source or in the target folder depends on the settings selected on the Missing Object tab of the *Options* dialog. A folder-to-folder source-to-target synchronization is one way of synchronizing child folders in one folder with those in another.
Bidirectional Sync	All the files in the source and target folders are synchronized. Merged files are checked into both views. Child folders that exist in both views may also be synchronized, depending on the selected options. What happens to objects that exist only in the source or in the target folder depends on the settings selected on the Missing Object tab of the *Options* dialog. A folder-to-folder bidirectional synchronization is one way of synchronizing child folders in one folder with those in another.

Table 8.1 Source Folder to Target Folder (continued)

Operation	Results
Missing Objects Only	View Manager ignores all objects that exist in both views. Thus, no merging occurs. What happens to objects that exist only in the source or in the target folder depends on the settings selected on the Missing Objects tab of the *Options* dialog. These objects can be ignored, added to the object from which they are missing, deleted from the object in which they exist, or both added and deleted (which results in a move). A folder-to-folder missing-objects-only synchronization allows you to ensure that two corresponding folders have all the same files without synchronizing any files.
Copy to Target	The source folder and its files are copied to the target folder as a child of that folder. The child folders of the source folder may also be copied, depending on the selected options. This operation also applies when the source object is a folder and the target object is a view. In this case, the source folder and its files are copied to the root folder of the target view as a child of that folder. The child folders of the source may also be copied, depending on the selected options.

Table 8.2 Source View to Target View

Operation	Results
Preview	The lower pane indicates the results of a source-to-target, bidirectional, or missing-objects-only synchronization, given the currently selected views and options. The preview starts with the two corresponding views and ignores the rest of the tree. Child views may be included, depending on the options selected. Child views follow the view root folder in the tree.
Source to Target Sync	All the folders and files in both views are synchronized. Merged files can be checked into the target. Child views in both views may also be synchronized, depending on the selected options. What happens to objects that exist only in the source or the target view depends on the settings selected on the Missing Objects tab of the *Options* dialog. A view-to-view source-to-target synchronization is one way to synchronize all child views in one view with all child views in another view.
Bidirectional Sync	All the folders and files in both source and target views are synchronized. Merged files can checked into both views. Child views in both views may also be synchronized, depending on the selected options. What happens to objects that exist only in the source or the target view depends on the settings selected on the Missing Objects tab of the *Options* dialog. A view-to-view bidirectional synchronization is one way to synchronize all child views in one view with all child views in another view.
Missing Objects Only	View Manager ignores all objects that exist in both views. Therefore, no merging occurs. What happens to objects that exist only in the source or the target view depends on the settings selected on the Missing Objects tab of the *Options* dialog. The objects can be ignored, added to the object from which they are missing, deleted from the object in which they exist, or both added and deleted (which results in a move). Views cannot be deleted (or moved) by View Manager, but only from StarTeam. A view-to-view missing-objects-only synchronization provides a way to copy all child views in one view to another view (if the Include Child Views option is selected). You can also use it to make sure that two corresponding views have the same files without synchronizing any files
Copy to Target	The source view is copied to the target view as a child of that view. Its child views may also be copied, depending on the selected options. This operation allows you to copy one view to another.

Using the Lower Pane

The lower pane displays the results of the operations you select.

Understanding the Tabs

The lower pane has two tabs, one for Preview and one for Run. The first tab is filled when you perform a preview operation. The second is filled when you perform a synchronization or copy-to-target operation.

Regardless of the tab, the pane contains a combination tree and list control. Each node of the tree has accompanying columns of information.

The pane is initially empty. If everything is in sync, a minimal tree contains only the root node (indicating the target object).

If you perform a preview operation followed by a run operation for the same source and target objects, View Manager displays the data from both trees so that you can compare them. Be aware that the contents of the source and target objects may change between the preview and the run if the server is not locked.

After a run operation completes, you can use the Run tab to:

- Determine what happened to your files. See "Finding Objects That Have Changed" on page 120.
- Compare or merge text files. See "Comparing and Resolving Text Files" on page 121.

Understanding the Tree

The tree shows the contents of both the source and the target objects. When the two objects do *not* have the same names, the name of the target object appears in the tree. The list also displays object properties in columns.

If you right-click a node in the tree and hold the mouse button down, the properties for the object associated with that node appear in a dialog. This is the same information that is displayed in columns in the lower pane, but it is presented vertically instead of horizontally.

To control the amount of data displayed in the tree, do one of the following:

- Select or clear the check mark in front of Display > Minimal Tree in Preview and Run Panes.

  When the check mark is selected, a minimal tree displays only projects, views, folders, and files that have changed or will change as a result of the run operation.

  When the check mark is cleared, View Manager displays the entire tree.
- Click the Show Complete Tree or Show Minimal Tree icon on the toolbar.

When you switch from a minimal to a complete tree (or vice versa), the lower pane initially displays a collapsed tree. You can expand or collapse branches of the tree as explained on page 93.

Understanding the Columns

The column headings for the lower pane are:

- Name, which lists the names of the source and target objects. When the names of these objects are different, the target object name is displayed in this column, which appears in both the Preview and Run panes.
- Locked By, which shows the name of the person locking the file and whether the file is locked in the source, target, or both. This column appears only in the Preview

pane. It alerts you to ask the user to unlock the file or to break the user's lock on the file (if you have the access rights to do so).

If either the source or target object is locked, the lock is reported, even if it may not cause the operation to fail. This saves the time it takes to examine all options and determine whether the lock will cause a problem. If both files are locked, this column contains: Source: <user name>; Target: <user name> ; if only one file is locked it contains only one of these.

If an operation fails due to a lock, lock information appears in the Result column in the Run pane.

- Preview Status or Run Status, a text description of the relationship between two corresponding objects. Every file has a status, as described in Table 8.3, "View Manager Statuses". The server configuration does not have a status. Projects, views, and folders have no status information unless they exist only in the source or the target object. This column appears in both the Preview and Run panes, but its name changes.
- Action, a text description of the action to be taken, given the status and option settings. This column appears in both the Preview and Run panes.
- Result, a text description of the outcome of the selected run operation on the two corresponding objects. This column displays only in the Run pane. All files have results, as described in the table. The server never shows a result. Projects, views, and folders show no results unless they are added, deleted, or moved.
- Source Revision, the number of the source tip revision.
- Target Revision, the number of the target tip revision.

Note The Source Revision and Target Revision columns appear in both the Preview and Run panes. Both tabs show the object's current tip revision numbers. The Preview tab displays the tip revision numbers before the run operation takes place; the Run tab shows the tip revision numbers after the run operation takes place.

Understanding Status, Actions, and Results

The following table lists View Manager statuses and explains the conditions that result in those statuses. For example, the run type affects the operations that are performed and, therefore, the resulting statuses.

Some View Manager statuses have names that are similar to the file statues used by StarTeam to indicate the relationship between a file's tip revision in the repository and the working file on your workstation. However, View Manager has nothing to do with working files. It compares tip revisions or corresponding files. These files may be in the same or different repositories.

If you have StarTeam open while you are performing an operation in View Manager, be aware that the status of your working files will change if their corresponding tip revisions change. For example, a working file's status changes from Current to Out Of

Date as View Manager checks a new tip revision into that repository. You may need to refresh the screen to see the working file's change in status.

Table 8.3 View Manager Statuses

Status	Action	Result
Current Corresponding files exist and have not changed since the last run operation or last recorded merge (if merge tracking is used).	*No action*	*No action taken*
Merge File has changed in both source and target.	*Manual merge* Resolution of conflicts or differences in text files will be manual and use the specified merge utility. For binary files, you may want to compare the two files in the application in which they were created.	*Manual merge needed*
	Automatic merge Automatic merge is possible because the corresponding text files have no conflicts.	*Merged* Merged file was checked into the target and/or source, depending on the operation and contents of the merged file.
	Use source file Source file is to be used for resolution of this text or binary file.	*Source file checked into target*
	Use target file Target file is to be used for resolution of this text or binary file.	*Target file checked into source*
Missing from Target Object exists in the source only.	*Add to target* Option settings indicate that object should be added.	*Added*
	Delete from source Option settings indicate that object should be deleted.	*Deleted*
	Move from source to target Option settings indicate that object should be deleted from the source and added to the target.	*Moved* Object was deleted from the source and added to the target.
	No action	*Shared from parent in source* Depending on view type, sometimes it inherits files from its new parent view.
Modified File has changed in target only.	*No action* Bilateral synchronization not specified.	*No action taken*
	Use target file Target file is to be used for resolution of this text or binary file in bidirectional synchronization.	*Target file checked into source*

Table 8.3 View Manager Statuses (continued)

Status	Action	Result
Not In Source Object exists in the target only.	*Add to source* Option settings indicate that this object should be added.	*Added*
	Delete from target Option settings indicate that this object should be deleted.	*Deleted*
	Move from target to source Option settings indicate that this object should be deleted from the target and added to the source.	*Moved* Object was deleted from the target and added to the source
	No action	*Shared from parent in target* Depending on the view's type, sometimes it inherits files from its new parent view.
Out Of Date File has changed in the source only.	*Use source file* A source-to-target synchronization checks the source file into the target.	*Source file checked into target*
	No action User specified a synchronization other than source-to-target.	*No action taken*
Exists In Both File exists in both the source and target but this is a missing-objects-only synchronization.	*No action* Missing-objects-only synchronization ignores files that exist in both the source and the target.	*No action taken*
Exists In Both Server configuration, project, view, or folder exists in both the source and target. Neither the option settings nor the run operation requires any action to be taken on the object itself.	*Check children* Object's children need to be checked for possible actions. For example, the object may have a child that needs to be added or shared.	*Checking done* Children have been checked for possible actions
Equivalent Name Objects with different names can be considered equivalent during a run operation. When the source object is the same type as the target object and a synchronization is selected as the run type, the objects are considered equivalent, even if their names are different. If the objects are views, their root folders are considered equivalent.	*Check children* The object's children need to be checked for possible actions. For example, the object may have a child that needs to be added or shared.	*Checking done* Children have been checked for possible actions

Table 8.3 View Manager Statuses (continued)

Status	Action	Result
Unknown Corresponding files exist in the source and the target but do not appear to be related. For example, unrelated files with the same name may have been added to both the source and the target since the last synchronization.	*No action* View Manager cannot determine what action should be performed.	*No action taken*
<any status>	*<any action to be taken>*	*Check-in failed* File could not be checked in. For example, if one or both of the corresponding files are locked by another user, you will see "source (or target) locked by <user name>"
Previously processed In a restarted operation, an object has this status if it was processed in a previous attempt to perform the operation.	*No action*	*No action taken*

Using the View Manager Graphical Interface

The following procedure provides a high-level overview of the View Manager graphical interface.

Overview

To use the View Manager graphical interface:

1. Start the application from the menu bar by selecting View > View Manager. After you launch View Manager, the StarTeam main window may retain focus. If this occurs, minimize the window.
2. Select source and target objects. See "Selecting Source and Target Objects" on page 101.
3. Select Tools > Options to set options for one or more types of run operations. Options can be set for the following operations:
   - Source-to-target
   - Bidirectional
   - Missing-objects-only
   - Copy-to-target

   These options affect:
   - Scope of the source and target objects. For example, if a folder is the source object, the scope indicates whether you are replicating only the selected folder or using the selected folder as the root of a branch to be replicated from the server tree shown in the upper left pane.

- Data that will be replicated. You may want all the historical revisions of each file or only its tip revision. You may want to copy the original date/time and user information for each replicated object. You may want to ignore missing files or add them.
- Handling of files that need to be merged.
- Locking of the server.

4 (Optional) Save the options to a properties file using File > Save.

In addition to the option settings for each type of run operation, the properties file stores information about the selected server configuration, the source object, and the target object. It does not store the password for logging onto the server nor the type of operation.

The properties file makes it easy and convenient to:

- Resynchronize every few days without having to reset all the options.

  Simply open the correct properties file by using File > Open, log onto the server, select a run operation, and View Manager does the rest.
- Copy one properties file and use it as the basis for another.

5 (Optional) Select a preview operation.

We recommend that you perform a preview operation before performing source-to-target, bidirectional, or missing-objects-only synchronizations, especially if you are unfamiliar with View Manager. This action enables you to see whether the operation will produce the results you expect. If it does not, you can adjust the option settings.

The copy-to-target operation does not have a preview operation because unexpected behavior is less likely. However, Borland recommends that you double-check the settings in the *Options* dialog before performing this operation.

After a preview operation, you can review files that need to be merged using a comparison utility, although you cannot merge files at this time. You can also locate locked files (if you are not breaking locks) and ask users to unlock them before a synchronization process begins.

6 Select a synchronization or copy-to-target operation.

7 After the operation completes, use Find Next Manual Merge and/or Find Previous Manual Merge to locate the files that need to be merged with a merge utility.

Note For more information about Visual Merge and Visual Diff, the default merge and comparison utilities, see the *StarTeam User's Guide*.

Selecting Source and Target Objects

After launching View Manager, you select the source and target objects by using the menu bar.

To select source and target objects:

1 Do one of the following:

- Select Server > Select Server from the menu bar.
- Click the Select Source Server icon on the toolbar.

The *Select a Server* dialog appears.

2 Select a server configuration from the Available Servers list box.

3 Click OK.

The server configuration appears in both panes.

4 Use the tree in the left pane to select a source object.

5 Use the tree in the right pane to select a target object.

Using a Properties File

Opening a properties file selects both a source and target object. The properties file contains information about the source and target objects and all the options used for the run—in short everything but the run type. A properties file has the extension .srm.

To open an existing properties file:

- Do one of the following:
  - Select File > Open from the menu bar.
  - Click the Open Properties File icon on the toolbar.

Note For more information on properties files, see “Using Properties Files” on page 118.

Configuring the Source View

To configure the source view to a particular view label or promotion state:

1 Select View > Configure Source View. The *Select a View Configuration* dialog appears.

2 Select the Labeled Configuration or Promotion State Configuration option button.

3 Select a view label or a promotion state from the drop-down list box below the selected option button.

4 Click OK.

Managing Server Configurations

To access a server configuration from View Manager, the server must appear in the servers list used by View Manager. To replicate data, you must have a user name, password, and the correct access rights.

Adding a Server Configuration

When you add a server configuration, you are really adding a server description to the starteam-servers.xml file used by StarTeam, StarTeam Server, and other StarTeam products.

To add a server:

1 From the main window, select Server > Select Server. The *Select a Server* dialog appears.

2 Click Add Server. The *Add Server* dialog appears.

See your administrator for the server address and endpoint information.

3 Enter an easy-to-remember description in the Server Description text box. This name should be unique. It is *not* case-sensitive and cannot contain colons (:).

4 Enter the address in the Server Address text box. The address is the name of the server or the IP address for the computer on which the server is running.

5 Enter the endpoint in the TCP/IP Endpoint text box. The endpoint is the port number.

6 Select the Compress Transferred Data check box if you want to use data compression.

7 Select an encryption type check box to encrypt data transferred between your workstation and the server. Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines.

   The encryption types are ordered (top to bottom) based on speed. Each encryption type is slower, but safer, than the type that precedes it.

   For more information about encrypting and compressing transferred data, see the *StarTeam User's Guide*.

8 Click OK.

Notes To change the server configuration properties, follow these same steps but click Server Properties instead of Add Server in step 2. To delete the server description, click Delete Server in step 2.

Locking a Server Configuration

If you will be performing more than one operation on a particular server, you may want to lock it manually to prevent others from checking in files and making other changes between your operations.

To lock a server:

- Select Server > Lock Server from the menu bar.

Performing Preview Operations

You can perform a preview only when objects of the same type have been selected from both the source and target panes. You must specify the type of synchronization desired (source-to-target, bidirectional, or missing-objects-only) and a set of options. View Manager then processes the selected objects and their children and ignores other branches of the tree.

A preview displays a composite tree in the lower pane when the Preview tab is selected. Statuses for each object in the tree indicate what would have happened if the operation had actually been performed. For example, if the preview compares two views and you note that only two branches of the folder hierarchy have changed since the last synchronization, you can change the source and target objects to the appropriate folders and synchronize just those branches.

After a preview, you can search the results to find differences, unresolved files, and locked files. If you double-click a pair of unresolved text files, the comparison utility enables you to compare files that have differences or conflicts. However, you cannot resolve differences or conflicts except during an actual run.

To run a preview operation:

1. From View Manager, select source and target objects. See “Selecting Source and Target Objects” on page 101.
2. Set options for the operation you want to perform. See “Setting and Storing Options” on page 108.
3. Do one of the following:
   - Select Run > Preview > Source to Target or Run > Preview > Bidirectional or Run > Preview > Missing Objects Only from the menu bar.
   - Click the Preview: Source-to-Target Sync , Preview: Bidirectional Sync , or Preview: Missing-Objects-Only Sync icon on the toolbar.
4. A confirmation dialog appears as a reminder that you that you should not run View Manager without current server backups. Although a preview operation makes no changes to either server, you should prepare for a subsequent run operation.
5. Click Review Options to check your option settings.
6. Click Run. The progress bars appear.
7. Assess what will happen by examining the data in the lower pane. You may want to reset some options to obtain different results.

   View Manager disconnects from the server containing the source and target objects only when you do one of the following:
   - Exit the application.
   - Open a new properties file.
   - Select a new source or target object.
8. Access the changes by examining the data in the lower pane. See “Finding Objects That Have Changed” on page 120. To compare a pair of text files that need resolution, select the files and double-click.

After you perform a Preview operation but before you start the actual synchronization or copy operation, it is advisable to search for locked files. After identifying the users who have locked these files, you can ask them to unlock the files or break their locks on the files. When all files are available, you can perform a successful synchronization or copy operation,

To find locked files, do one of the following:

- Select Find > Next Locked File or Find > Previous Locked File from the menu bar.
- Click the Find Next Locked File icon or Find Previous Locked File icon on the toolbar.

Performing Synchronization Operations

View Manager allows you to perform three different types of synchronization:

- Synchronizations that merge corresponding files.
  - Source-to-target synchronization can be performed after objects of the same type have been selected from the source and target panes. It merges files and checks them into the target only. To synchronize objects from target to source, the contents of the source and target panes must be reversed.
  - Bidirectional synchronization can be performed after objects of the same type have been selected from the source and target panes. It merges files and checks them into both the source and target.

    For objects that exist only in the source or only in the target, both types of synchronization allow objects that have no corresponding objects to be ignored,

added, or deleted, depending on the settings for missing objects in the Options dialog. Folders and files can be moved that is, added to the new location and deleted from the old location.

- Missing-objects-only synchronization. This type of synchronization ignores the objects that are found in both the source and target objects. Therefore, it merges no files. It ignores, adds, deletes, or moves objects that exist only in the source or only in the target (depending on the selected options).

To synchronize objects from target to source, the contents of the source and target panes must be reversed.

Synchronization does *not* mean that the source and target objects become identical afterward. For example:

- During source-to-target synchronizations, merged files are checked in to the target only.
- The options set for projects, views, folders and files that exist in only the source or only the target can keep the source and target objects from becoming identical.

To run a synchronization operation:

1. Make a complete backup of the source and target server configuration databases, repositories, and so on. See Appendix C, “Backing Up the Application,” on page 249.
2. From View Manager, select source and target objects. See “Selecting Source and Target Objects” on page 101.
3. Set options for the operation you want to perform. See “Setting and Storing Options” on page 108.
4. Do one of the following:
   - Select Run > Sync > Source to Target or Run > Sync > Bidirectional or Run > Sync > Missing Objects Only from the menu bar.
   - Click the Source-to-Target Sync, Bidirectional Sync, or Missing-Objects-Only Sync icon on the toolbar.
5. A confirmation dialog appears as a reminder that you that you should not run View Manager without current server backups.
6. Click Review Options to ensure that your option settings are correct. Although you can cancel an operation, changes occurred before the cancellation cannot be undone.
7. If you have the appropriate backups and option settings, click Run. The progress bars appear.
8. After the operation is complete, access the changes by examining the data in the lower pane. See “Finding Objects That Have Changed” on page 120.

If you have performed a source-to-target or bidirectional synchronization, some text files may require manual merging. See “Comparing and Resolving Text Files” on page 121.

Performing Copy-to-Target Operations

You can copy a source object in entirety to a target object by selecting an object on the right pane that is *one* type higher than the source (in the left pane). The hierarchy from top to bottom is server/project/view/folder.

If a view is selected in both the source and the target, the source view can be copied to the target view as a child of the target view. If a folder is selected in both the source and the target, the source folder can be copied to the target folder as a child of the target folder.

After the copy operation completes, the copied object has the same working folders in both the source and the target locations. It may be desirable to change the working folders for one of the objects.

To include child views with the copied view or include child folders with the copied folder, you must select the Include Child Views or Include Folder Views check box, respectively, on the General tab of the *Options* dialog.

To run a copy-to-target operation:

1. Back up the target server configuration database, repository, and so on. See Appendix C, “Backing Up the Application,” on page 249 for backup information.
2. From View Manager, select source and target objects. See “Selecting Source and Target Objects” on page 101.
3. Set options for the operation you want to perform. See “Setting and Storing Options” on page 108.
4. Do one of the following:
   - Select Run > Copy to Target from the menu bar.
   - Click the Copy-to-Target icon on the toolbar.
5. A confirmation dialog appears. It reminds you that you should not be running View Manager without current server backups. The target server is the only server that is changed by a copy-to-target operation.
6. Click Review Options to review your option settings.

   View Manager operations can take a long time. You want to be certain that you have selected the correct options. You can cancel the operation, but changes that happened before you cancelled cannot be undone.
7. Click Run.

   The progress bars appear.
8. Assess what happened by examining the data in the lower pane.

Important View Manager does not disconnect from the server containing the source and target objects until you exit the application, open a new properties file, or select a new source or target object.

If you disconnect prior to resolving all the files, you can check the log for the names of the files or restart View Manager and redo the run.

Observing Progress

While performing a run operation, View Manager displays up to three progress bars, each showing a percentage of completion.

- The first bar displays over-all progress and is always displayed.
- The second bar displays project level progress. It appears when the source object is a project or server.
- The third bar displays folder-level progress and is always displayed.

Below the progress bars are statistics about the total number of projects, views, and folders that have been processed.

You can click the Cancel button on the progress dialog to stop the run operation.

Resuming an Aborted Operation

If View Manager failed to complete an operation,.any part of the operation that was completed cannot be automatically undone.

You can resume an operation if:

- View Manager stopped the operation because its retries were exhausted.

  The application enables you to set a number of retries and the interval between them. The default setting is three retries at one-minute intervals. If View Manager successfully restarts during the retry cycle, the retry counter is reset to 0.

  If the retries all fail, you are prompted to save the source and target selections and your current option settings to a properties file.

- You cancelled the operation.
- View Manager stopped the operation because a fatal exception occurred.

  In this situation, View Manager resumes with the same source and target objects and option settings that it had previously. It loses the tree that was built in the first attempt, but continues from the point at which it left off. If some files in the lost section of the tree required manual merging in the lost section of the tree, you can find entries about them in the log. The entries will look similar to these for viewlog.bmp:

  ```
  9/4/04 8:55:24 AM PDT  Examining source File StarDraw\StarDraw\StarDraw\
  viewlog.bmp:1329
  9/4/04 8:55:24 AM PDT  File found in target folder
  9/4/04 8:55:24 AM PDT  Synchronizing File StarDraw\StarDraw\StarDraw\
  viewlog.bmp:1329
  9/4/04 8:55:24 AM PDT  Synchronizing item StarDraw\StarDraw\StarDraw\
  viewlog.bmp:1329
  9/4/04 8:55:24 AM PDT  Both items have changed since the common ancestor
  9/4/048:55:24AMPDTFileStarDraw\StarDraw\StarDraw\viewlog.bmp:1329isin
  conflict.
  ```

To restart an operation from the point at which it left off:

- Select Run > Resume from the menu bar.

Using the Log Files

The log files for each run operation have the extension .log. Each log file is identified by its creation data and time, which is included in its name. For example, a .log file created on September 4, 2003 at 9:03:02 A.M. would be named Repmgr-04-Sep-03-09-03-02.log.

When a log file contains approximately 5MB of data, a new log file is created. The log files are stored in the Log folder, which is a child of the installation folder. For example, if you installed View Manager with the Server, this folder is usually C:\Program Files\Borland\StarTeam Server x.x\Log.

The .log files are accompanied by tree.txt files, which list the information that would appear in the lower pane's tree. For example, Repmgr-04-Sep-01-09-03-02.log is accompanied by Repmgr-04-Sep-01-09-03-02-tree.txt.

The first few lines of the log identify the product, its copyright date, and the StarTeam Runtime build number.

- To locate problems, search for "Exception thrown".
- To find out what happened to a specific object, search for that object by name, for example, viewlog.bmp.

- To find out more about newly added objects, search for “Created project”, “Created view”, “Created folder”, “Created file”, or just “Created”.

To view the log from View Manager

1. Select Run > View Log from the menu bar. The *Select a Log File* dialog appears.
2. Select a log file.
3. Click OK. The *View Log* dialog appears.
4. After reviewing the log, do one of the following:
   - Click Close to exit.
   - Click Browse to select another log file.

Understanding Restart Files

Restart files have the extension .rsm. Each .rsm file indicates the point at which a run operation stopped before it abnormally terminated. The file is used to resume the operation and is deleted by View Manager when the operation is successful.

The .rsm file lists the objects that have been completely processed and will not be re-examined during when the operation is resumed. If all the files and child folders in a folder have been completely processed, then only the folder name is listed, not the names of the objects that the folder contains. Similarly, after a view or project has been completely processed, only the view or project name is listed.

The .rsm files are stored in the folder in which streplicate.jar was installed. If you wish to resume an operation from the command line, you must provide the path to the correct .rsm file for the operation.

Setting and Storing Options

View Manager provides an *Options* dialog that enables you to select options for different types of run operations and save them to a properties file for later reuse. Not all View Manager options apply to all types of run operations.

The *Options* dialog has a drop-down list box from which you can select:

- Source to Target
- Bidirectional
- Missing Objects Only
- Copy to Target

Each item in the list box represents a type of run operation. When you select a specific type of run operation, the tabs and options available for that operation appear.

All run operations share at least some of the options. For example, the General and Locking tabs apply to all operations. When options apply to more than one type of run operation, you may want to select different settings, depending on the type of operation.

Some option settings for a particular run operation will be ignored if they do not apply to the selected source and target objects. For example, the Include Child Folders option is ignored unless you select a folder as the source object.

Setting General Options

Use the General tab of the *Options* dialog to:

- Set the scope of the replication when the source object is a view or folder.

- Include child views

  Indicates how much of the tree to include in the run operation. This setting is valid only when a view is selected as the source object and the view configuration is current. When you are synchronizing views, it applies to both the selected source and target views.

  This option is always enabled. The default is cleared.

- Include child folders
- Indicates how much of the tree to include in the run operation. This setting is valid only when a folder is selected as the source object. When you are synchronizing folders, it applies to both the selected source and target folders.

  This option is always enabled. The default is selected.

- Specify the extra information that should be replicated.
  - (Replicate) Users

    Enables View Manager to use the original user names in the “Created By” properties of created objects.

    This option is always enabled. The default is cleared.

  - (Replicate) Creation times

    Copies creation times, along with created items. When cleared, the object’s creation times are the times at which View Manager has created them.

    If the source and target objects are the same view or folders in the same view, objects will be shared rather than added. Therefore, these objects will have the original creation times, regardless of the setting for this option.

    This option is always enabled. The default is cleared.

  - (Replicate) View labels

    Copies all the view label definitions from one view to the other. Depending on the missing-object option settings and the type of operation, view label definitions can be copied from the source to the target and the target to the source.

    This option enables View Manager to include the view labels attached to file revisions created as part of the run operation.

    View labels are valid at the view level. For example, if you select a folder as the source object, view label information from the view containing that folder is copied to the target view in which the folder is replicated.

    If the source and target objects are the same view or in the same view, this option. is not needed.

    This option is disabled unless Creation Times is selected. The default is cleared.

Note View labels and promotion states that have the same names in both the source and target view are never copied. Thus. if Build 1 is a view label in both locations, that label can never be changed in either location by View Manager.

If the source and target objects are the same view or folders in the same view, this option is not needed.

To avoid problems, you should maintain view labels in only one of the two locations. Remember that the application can change the name of a view label, if necessary.

  - (Replicate) Promotion states

    Copies all promotion state definitions from one view to the other. Depending upon the missing-object option settings and the type of operation, promotion state definitions can be copied from the source to the target and the target to the source.

Promotion states are valid at the view level. For example, if you select a folder as the source object, the promotion state information from the view containing that folder is copied to the target view in which the folder is replicated.

If the source and target objects are the same view or folders in the same view, this option is not necessary.

This option is disabled unless View Labels is selected. The default is cleared.

- (Replicate) Revision labels

  Copies revision labels on an as-needed basis, as file revisions are created.

  In other words, if a file revision that is being copied has revision labels, the revision labels will be attached to the copy as well as to the original file. If the label does not already exist in both source and target, View Manager creates it.

  If the source and target objects are the same view or in the same view, this option is not necessary.

  This option is always enabled. The default is cleared.

- Control the number and frequency of retries for problematic server connections.
  - Number of retries

    Specifies the number of attempts to reconnect that View Manager will make to a given server.

  This option is always enabled. The default is 3 times.

  - Time between retries ___ minutes

    Specifies the number of minutes between reconnection attempts.

    This option is always enabled. The default is one minute, but it may take a few minutes for StarTeam Runtime to detect the disconnect.

Setting Missing Objects Options

The Missing Objects options are identical in source-to-target, bidirectional, and missing-objects-only operations. These options do not apply to copy-to-target operations because, in a copy-to-target operation, everything is missing from the target. The options on this tab apply to objects missing from the source or target objects within the scope of the operation.The options on this tab apply to objects missing from the source or target objects within the scope of the operation.

In the following figure, the scope includes all views, folders, and items in the two projects. Project A and Project B are treated as equivalent. Project A has a child view that Project B does not have. During synchronization of Project A and Project B,

Project A's child view would be added to Project B if the Add Views To Target check box is selected.

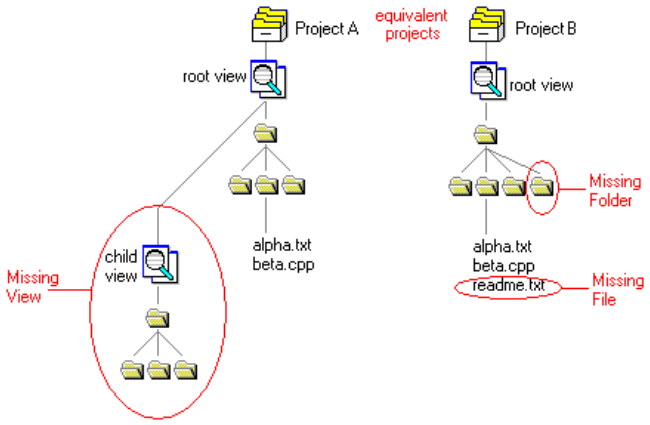

Similarly, the folder in Project B that is labeled as "Missing Folder" would be added to Project A if the Add Folders To Source check box is selected, and the file labeled "Missing File" would be added to Project A if the Add Files To Source check box is selected.

Use the Missing Objects tab to do the following:

- Only in the source:
  - Add projects to the target

    This option is always disabled.
  - Add views to the target

    Allows views that exist only in the source object to be added to the target object. Such views are copied in entirety to the target project, including the view, its child views, and their folders and files.

    No view is added based on this option unless its parent (or an equivalent object) exists in both the source and target objects.

    This option applies when a view is the source object only if the Include Child Views option is also selected. It does not apply when a folder is the source object.

    If a view is the selected source object, that view is *not* missing. If a view is the root view of the selected project, it is also *not* missing.

    This option is always enabled. The default is selected.
  - Add folders to the target

    Enables folders that exist only in the source object to be shared to the target object. Such folders are shared in entirety to the target object, including the folder and its child folders and files, recursively.

    No folder is shared based on this option unless its parent (or an equivalent object) exists in both the source and target objects.

    This option applies when a view is the source object. It applies when a folder is the source object only if the Include Child Folder option is also selected.

    If a folder is the selected source object, that folder is *not* missing. If a folder is the root folder of the selected view or the root folder of the root view of the selected project, it is also *not* missing.

    This option is always enabled. The default is selected.
  - Add files to the target

Enables files that exist only in the source object to be shared to the target object.

No file is shared based on this option unless its folder (or an equivalent folder) exists in both the source and target objects.

This option is always enabled. The default is selected.

- Delete folders from the source

  Enables folders that exist only in the source object to be deleted from the source object. Such folders are deleted in entirety from the source object, including the folder and its child folders and files.

  No folder is deleted based on this option unless its parent (or an equivalent object) exists in both the source and target objects.

  This option applies when a view is the source object. It applies when a folder is the source object only if the Include Child Folder option is also selected.

  If a folder is the selected source object, that folder is *not* missing. If a folder is the root folder of the selected view or the root folder of the root view of the selected project, it is also *not* missing.

  This option is always enabled. The default is cleared.

- Delete files from the source

  Enables files that exist only in the source object to be deleted from the source object.

  No file is deleted based on this option unless its folder (or an equivalent folder) that exists in both the source and the target.

  This option is always enabled. The default is cleared.

- Only in the target:
  - Add projects to the source

    This option is always disabled.

  - Add views to the source

    Enables views that exist only in the target object to be added to the source object. Such views are copied in entirety to the source project, including the view, its child views, and their folders and files.

    No view is added based on this option unless its parent (or an equivalent object) exists in both the source and target objects.

    This option applies when a view is the target object only if the Include Child Views option is also selected. It does not apply when a folder is the target object.

    If a view is the selected target object, that view is *not* missing. If a view is the root view of the selected project, it is also *not* missing.

    This option is always enabled. The default is cleared.

  - Add folders to the source

    Enables folders that exist only in the target object to be shared to the source object. Such folders are shared in entirety to the source object, including the folder and its child folders and files.

    No folder is shared based on this option unless its parent (or an equivalent object) exists in both the source and target objects.

    This option applies when a server, project, or view is the target object. It applies when a folder is the target object only if the Include Child Folder option is also selected.

If a folder is the selected target object, that folder is *not* missing. If a folder is the root folder of the selected view or the root folder of the root view of the selected project, it is also *not* missing.

This option is always enabled. The default is cleared.

- Add files to the source

  Enables files that exist only in the target object to be shared to the source object.

  No file is shared based on this option unless its folder (or an equivalent folder) exists in both the source and target objects.

  This option is always enabled. The default is cleared.

- Delete folders from the target

  Enables folders that exist only in the target object to be deleted from the target object. Such folders are deleted in entirety from the target object, including the folder and its child folders and files.

  This option applies when a server, project, or view is the target object. It applies when a folder is the target object only if the Include Child Folder option is also selected.

  No folder is deleted based on this option unless its parent (or an equivalent object) exists in both the source and target objects.

  If a folder is the selected target object, that folder is *not* missing. If a folder is the root folder of the selected view or the root folder of the root view of the selected project, it is also *not* missing.

  This option is always enabled. The default is cleared.

- Delete files from the target

  Enables the deletion of files from the target object that exist only in the target object.

  No file is added based on this option unless its folder (or an equivalent folder) exists in both the source and target objects.

  This option is always enabled. The default is cleared.

Note To delete projects and views from the source or target, use StarTeam. To prevent accidental loss of data, View Manager does *not* provide delete project and delete view options. If you accidentally delete folders or files while using View Manager, you can roll back to the appropriate view in StarTeam and recover those folders and files.

Setting Merge Options

View Manager can track merge operations for you, if you record them. This feature is especially convenient if you merge the same two views more than once in the same direction. With merge tracking, View Manager remembers the last time you merged two corresponding files (one from the source and one from the target view). Regardless of the changes that were included or excluded during the merge process, View Manager considers the two files to be identical (in the direction of the merge) after the merge. Thus, View Manager reports that the two corresponding files need to be merged only when new changes are made to either or both of the files.

When the merge tracking feature is turned off, the application considers any two corresponding files to be different unless they have exactly the same contents.

The Merge options are valid only if the run operation is source-to-target or bidirectional synchronization. In missing-objects-only or copy-to-target operations, nothing is merged. Therefore, the Merge tab is not included on the *Options* dialog for missing-objects-only or copy-to-target operations.

Use the Merge tab to:

- Merge recording/tracking
  - Record each merge

    Records the revisions involved in a specific merge. Recording and tracking merges goes together.

    This option is always enabled. The default is cleared.
  - Use merge tracking for merge resolution

    Uses the most recent recorded merge as a basis for the status of a given file during a preview or for the actual merge operation. For example, if neither of the corresponding files has changed since the last merge, the status is Unchanged. If merge tracking is not selected, corresponding files are merged if they are not identical. If merge tracking is selected, the files can be different and still not require a merge operation.

    This option is always enabled. The default is cleared.
    - Use source as common ancestor

      Bases the merge on the source file from the last recorded merge as well on the pair of files being merged.

      If merge tracking is not selected, the most recent common ancestor is used to perform the three-way merge. The common ancestor is a historical revision with the same revision number in both views.

      When the two files are in different server configurations, View Manager iterates through the history of corresponding files and uses the MD5 checksum to identify a common ancestor—that is, a pair of “matching” revisions, one from the source and one from the target. If no common ancestor is found, the pair of files has the status Unknown.

      This option is enabled when the Use merge tracking check box is selected. The default is selected.
    - Use target as common ancestor

      Bases the merge on the target file from the last recorded merge as well as on the pair of files being merged.

      If merge tracking is not selected, the most recent common ancestor is used to perform the three-way merge. The common ancestor is a historical revision with the same revision number in both views.

      When the two files are in different server configurations, View Manager iterates through the history of corresponding files and uses the MD5 checksum to identify a common ancestor—that is, a pair of “matching” revisions, one from the source and one from the target. If no common ancestor is found, the pair of files has the status Unknown.

      This option is enabled when the Use merge tracking check box is selected. The default is cleared.

Setting Utility Options

The options on the Utilities tab are valid only if the run operation is a source-to-target or bidirectional synchronization.

The default utilities are Visual Merge and Visual Diff. If you prefer, you can select alternate utilities and designate options to be used with them. If you designate the options incorrectly, your utility may start up with no files or the wrong file may be used as the ancestor file in a three-way merge.

- Visual Merge

  Designates Visual Merge as your merge utility on Windows platforms.

This option is always enabled. The default is selected.

- Ignore whitespace

  When you compare two lines of text files, this option ignores trailing white space and treats all other strings of white space as equal in length. For example, the following lines are equivalent:

  " hi mom "
  " hi mom"

  This option is enabled only if the Visual Merge option button has been selected. The default is cleared.

- Alternate merge utility

  Enables you to specify a merge utility other than Visual Merge.

  This option is always enabled. The default is cleared.

  - Merge utility

    Path to the alternate merge utility.

  This option is enabled only if the Alternate merge utility option button has been selected. The default is blank.

  - Options

    Lists the command-line options that go with the alternate merge utility. Use the following terms to represent files sent by View Manager to the merge utility:

    $branchtip A place holder for the path to the tip revision of the file to be merged.

    $usertip A place holder for the path to the local working file to be merged.

    $basefile A place holder for the path to the common ancestor for the $branchtip and $usertip files.

    $resultfile A place holder for the path to the file that will store the output from the merged file.

    For example, suppose the Merge Utility text box contains "D:\Programs\Merge Utility 5\Mergeutil.exe" and the Options text box contains `"-s $branchtip $usertip $basefile $resultfile"`. View Manager assumes that the options entered are appropriate for the utility to be used. Each time that it needs to call the merge utility, it replaces the terms that start with $ with the actual paths to those files and asks the operating system to execute the completed command.

    This option is enabled only if the Alternate merge utility option button has been selected. The default is blank.

- Alternate comparison utility

  Enables you to specify a merge utility other than Visual Diff.

  This option is always enabled. The default is cleared.

  - Comparison utility

    Path to the alternate comparison utility.

    This option is enabled only if the Alternate comparison utility check box has been selected. The default is blank.

  - Options

    Lists the command-line options that go with the alternate comparison utility. Use the following terms to represent files sent by View Manager to the merge utility:

    $file1 A place holder for the path to the first of the two files to be compared.

$file2 A place holder for the path to the second of the two files to be compared.

For example, suppose the comparison utility's path is "D:\Programs\Compare Utility 8\Diffutil.exe" and the options to be used are `"$file1 $file2"`. View Manager assumes that the options are appropriate for the utility to be used. Each time it needs to call the comparison utility, it replaces the terms that start with $ with the actual paths to those files and asks the operating system to execute the completed command.

This option is enabled only if the Alternate comparison utility check box has been selected. The default is blank.

Setting Locking Options

The Locking tab allows you to lock the server exclusively during a run or preview operation, preventing users from changing the contents of one or both repositories during your operation.

Note You can also use a menu command (Server > Lock Server) to control the locking and unlocking of the server.

Use the Locking tab:

- On Run
  - Exclusively lock server on run

    Exclusively locks the server when a run operation begins.

    This option is always enabled. The default is cleared.

    - Unlock server after run

      Unlocks the server when a run operation ends. The operation ends when the last two files have been examined and reported on. A message appears explaining whether or not the run was successful.

      This option is disabled unless the Exclusively Lock Server On Run check box is selected. The default is cleared.
- On Preview
  - Exclusively lock server on preview

    Exclusively locks the server when a preview operation begins. This option is always enabled. The default is cleared.
  - Unlock server after preview

    Unlocks the server when a preview operation ends. The operation ends when the last two files have been examined and reported on. A message appears explaining whether or not the run was successful.

    This option is disabled unless the Exclusively Lock Server On Preview check box is selected. The default is cleared.

Setting File Options

Use the File tab for the following options:

- All files
  - Copy the history

    Copies the file history (and the appropriate vault files) to the source and/or the target, depending on the direction of the operation. If not selected, only the tip revisions of each file will be used.

This option does *not* apply to files that have changed in both the source and target. If the operation is missing-objects only or copy-to-target, this option is the only one that appears on the File tab, because no merging is possible.

If one file is out of date with respect to the other, then all newer versions of the current file are copied, to bring the old file up to date. If both files have changed, then this action is impossible, so View Manager checks in only the merged file.

Unless you select both Copy File History and the (Replicate) Creation Times check box on the General tab, the revision times are the times at which View Manager created them during this operation. If the (Replicate) Creation times check box has been selected, View Manager copies the original creation times to the new location.

This option is always enabled. The default is cleared.

- Break lock if file needs to be checked in:

  When selected, enables View Manager to check in files even if they are locked by someone else. You must have the access rights necessary to break locks.

  Valid only for source-to-target and bidirectional synchronizations.

  This option is always enabled. The default is cleared.

- Resolve differences for text files

  Files conflicts or differences can be resolved automatically by using the merge utility or by substituting the existing source or target file for the merged file.

  - Merge automatically

    Determines whether all text files or only text files with conflicts will require manual merging with a merge utility.

    When selected, View Manager uses the StarTeam Runtime's merge command to merge text files that have no conflicts. The remaining File option settings indicate how to resolve conflicts.

    When cleared, View Manager does no automatic merging of text files. The rest of the File option settings indicate how to resolve the differences in any text file with Merge status.

    This option is always enabled. The default is cleared.

  - For all merging

    - Use specified merge utility

      Enables conflicts to be resolved with the merge utility.

      This option is always enabled. The default is selected.

    - Use existing file as merged file

      Enables View Manager to use either the source or target file as the merged file.

      This option is always enabled. The default is cleared.

      - Use source file. Enabled when Use existing file as merged file is enabled. The default is selected.
      - Use target file. Enabled when Use existing file as merged file is enabled. The default is cleared.

- Resolve differences for binary files

  View Manager cannot merge binary files. Binary files with differences can either be ignored or resolved by substituting the existing source or target file for the merged file.

  - Use existing file as merged file

Enables View Manager to use either the source or target file as the merged file.

When cleared, binary files are not merged. The pair has the status Merge, the Action column says “Manual merge.”

This option is always enabled. The default is selected.

- Use source file. Enabled when Use existing file as merged file is enabled. The default is selected.
- Use target file. Enabled when Use existing file as merged file is enabled. The default is cleared.

- Lock unresolved files

  Enables you to lock files that have not been resolved. Files are unresolved if they have Merge status but cannot be resolved automatically. If they are text files, they either have conflicts or differences, so merging cannot be done automatically. If they are binary files, neither the source or the target file has been selected as the merged file.

  - Text

    Controls the locking of unresolved text files.

    This option is always enabled. The default is cleared.

  - Binary

    Controls the locking of unresolved binary files.

    This option is always enabled. The default is cleared.

Using Properties Files

In View Manager, a properties file contains information about the selected server, configuration, the source and target objects, all options used for the run, the name of the log file, and the name of the restart file. It does not store the password for logging onto the server nor the type of operation. A properties file has the extension .srm.

Once you have created a properties file, it can be opened any time you want to reuse the same source and target objects, regardless of the run operation to be performed at that time. It makes it easy to resynchronize every few days without having to reset all the options.

You can open the file from inside View Manager or start View Manager with a specific properties file by double-clicking the file in Windows Explorer.

To create a properties file:

1. Select source and target objects. See “Selecting Source and Target Objects” on page 101.
2. Select appropriate options. See“Setting and Storing Options” on page 108.
3. Do one of the following:
   - Select File > Save from the menu bar.
   - Click the Save Properties File icon on the toolbar.

   The *Save Properties File* dialog appears.
4. (Optional) Click Change Options to review or change any of the options to be stored with this file. The *Options* dialog appears.
5. When you are satisfied that the options are correct, click OK to return to the *Save Properties File* dialog.
6. Click OK.

Tip To create a properties file with settings that are similar to those of an existing file, open the existing file, make changes, and save it with a new name.

The following example shows properties and some possible option settings. The comments (//) have been added to clarify what each property represents. Actual properties files are uncommented.

```
//identify source computer, user, and object
//in this case, a folder in the StarDraw project's
//root view
sourceaddress=localhost //could have been IP address
sourceendpoint=3000
sourceusername=Administrator
sourcefolder=StarDraw\Source Code\
sourceproject=StarDraw
sourceview=StarDraw

//identify target computer, user, and object
targetaddress=bonzai //could have been IP address
targetendpoint=49201
targetusername=Administrator
targetfolder=Big Project\Source Code\
targetproject=Big Project
targetview=Big Project

//all the ".0" options are source-to-target options
arr_syncflags.0=117489727
arr_mergeflags.0=69664
arr_retryattempts.0=3
arr_retrywait.0=1

//all the ".1" options are bidirectional options
arr_syncflags.1=1022
arr_mergeflags.1=69664
arr_retryattempts.1=3
arr_retrywait.1=1

//all the ".2" options are missing-objects-only options
arr_syncflags.2=62
arr_mergeflags.2=65568
arr_retryattempts.2=3
arr_retrywait.2=1

//all the ".3" options are copy-to-target options
arr_syncflags.3=62
arr_mergeflags.3=65568
arr_retryattempts.3=3
arr_retrywait.3=1

//paths from server description to folder
SOURCE_Path=blackhole;StarDraw;StarDraw;StarDraw;Source Code
TARGET_Path=orion;Big Project;Big Project;Big Project;Source Code

//computer names
SOURCE_ServerHost=localhost
TARGET_ServerHost=localhost

//port numbers
SOURCE_ServerPort=3000
TARGET_ServerPort=3000
```

```
//protocol
int_SOURCE_ServerProtocol=6
int_TARGET_ServerProtocol=6

//server name
SOURCE_ServerName=orion
TARGET_ServerName=orion

//server descriptions
TARGET_ServerDescription=orion
SOURCE_ServerDescription=orion

//compression level
int_TARGET_ServerCompression=0
int_SOURCE_ServerCompression=0

//encryption level
SOURCE_ServerEncryption=NULL
TARGET_ServerEncryption=NULL

//lock status
int_SOURCE_Lock=0
int_TARGET_Lock=0
```

Understanding the Configuration File

The View Manager file that stores the current values of commands on the Display menu is named sbrmpref.ini. For example, if you show the complete tree in the lower pane, the sbrmpref.ini file contains the following line:

```
ShowFullConflictTree=true
```

Finding Objects That Have Changed

Differences are determined by status and apply to all objects other than the server, which never has a status. See Table 8.3, “View Manager Statuses,” on page 98 for more details.

All statuses, with the exception of Current, represent differences.

The icon for unresolved files displays for any pair of corresponding text files that require merging. These files have the result state “Manual merge required.”

When text files are merged automatically, unresolved text files contain conflicts. (A conflict is defined as a difference between two corresponding lines of two corresponding text files.) When text files are not merged automatically or by substitution of the source or target file, unresolved text files can contain any difference. View Manager contains a merge utility so that you can resolve conflicts or differences.

For binary files, a conflict is defined simply as a difference between the two corresponding files. Unless you substitute the source or target file as the merged file, all binary files are unresolved.

In all cases, information about unresolved files is logged.

By using the commands on the File menu, you can locate files and other objects that have changed.

To find files that need to be manually merged, do one of the following:

- Select Find > Next Manual Merge or Find > Previous Manual Merge from the menu bar.

- Click the Find Next Manual Merge or Find Previous Manual Merge icon on the toolbar.

To find files with differences, do one of the following:

- Select Find > Next Difference or Find > Previous Difference from the menu bar.
- Click the Find Next Difference or Find Previous Difference icon on the toolbar.

Comparing and Resolving Text Files

Unresolved files can be located by using the commands on the Find menu or their corresponding toolbar buttons.

After a preview operation, double-clicking text file names in the lower pane displays the files in the comparison utility.

After a synchronization operation, double-clicking the text file names displays the merge utility.

To compare or resolve text files:

1. Use the Next Manual Merge and Previous Manual Merge menu commands or toolbar buttons to locate a set of unresolved corresponding files.
2. Double-click the name of a file to display the merge utility.
3. Do one of the following:
   - Use the comparison utility to review these files. See the *StarTeam User's Guide* for details on using Visual Diff.
   - Use the merge utility to resolve the conflicts or differences between these files. See the *StarTeam User's Guide* for details on using Visual Merge.

Note Be aware that View Manager does *not* update the revision numbers in the lower pane when merged files are checked in.

Resolving Binary Files

Unresolved binary files can be located by using the commands on the Find menu or their corresponding toolbar buttons. They, too, are flagged, but double-clicking them does not display the merge utility. (In fact, it displays a "Merge operation failed" message because View Manager cannot merge binary files.)

To compare binary files and resolve their differences, you can often use the applications with which they files were created. For example, Microsoft Word can be used to compare and merge .doc or .rtf files.

Identifying Files Checked in by View Manager

The application can be used to identify files that were replicated by View Manager. These files have RM as the first two characters in the Comment field, followed by the date and time at which View Manager checked in the file. The data and time are followed by the original comment provided for the file.

When a merged file is checked in as a new revision, the revision's Comment field provides merge information. The comment identifies the source item that was used in the merge when a file was checked into the target and vice versa. For example, if you merge files, the comment might be "RM <date/time>: Merge from MAINFRM.CPP, Revision 1.3.1.9".

To find these files using StarTeam:

1. Start the application and open the appropriate project view.

2. Display the Comment field in the upper pane.
   a. Do one of the following:
      - Right-click a column header on the upper pane, then select Show Fields from the context menu.
      - Select Filters > Show Fields from an appropriate menu or context menus.

      The *Show Fields* dialog displays two lists. The Available Fields list box contains all fields that could be displayed as column headers but are not currently displayed. The Show These Fields In This Order list box displays all fields that are currently displayed.
   b. Double-click the Comment field in the Available Fields list box to move it to the Show These Fields In This Order list box.
   c. Click OK.

   To change the order of the fields to be displayed in the upper pane, drag each field name to the desired location in the Show These Fields In This Order list box.
3. Find files that have “RM” contained in the Comment field of their tip revisions.
   a. Select File > Find from the menu bar. The *Find* dialog appears.
   b. Enter RM in the Search For text box.
   c. From the Look In group box, select the This Field option button.
   d. Select Comment from the drop-down list box.
   e. Select the Match Case check box.
   f. Click OK to locate the first file.
4. Press *F3* and *Shift+F3* to find the next or previous files as appropriate.

Tip You can also sort files based on the contents of the Comment field in their tip revisions. See the *StarTeam User’s Guide* for details.

Using the View Manager Command Line

From the Cross-Platform client, you can also launch and use View Manager View Manager from the command line. Although you start the command-line differently on different platforms, the same options apply.

By default, the View Manager command line:

- Performs a source-to-target synchronization
- Performs an automatic merge of corresponding text files that have changed in both the source and target
- Displays Visual Merge whenever a pair of corresponding text files contain conflicts. You can use the merge utility to resolve conflicts before View Manager checks in the merged file (if appropriate and possible) and goes on to the next pair of files. Checking in is not appropriate if the merged file is identical to the current tip revision. It is not possible if you do not have the access rights allowing you to check in files or break another user's lock on the file.

Starting View Manager from the Command Line

To start View Manager from the command-line interface, from any platform, use a command similar to the following, followed by the options you intend to use:

```
java-Xmx49201m-classpath"streplicate.jar:jhall.jar:starteam-sdk.jar"
com.starbase.starteam.streplicate.streplicator
```

Depending on your system setup and the location from which you run the command, you may need to replace java with the complete path to the JRE.

Important

We recommend using Java's -Xmx*n* option in the command that for the command-line interface. Borland suggests replacing *n* with a number that represents 80% of the total memory available on your computer.

Here is the documentation that Sun currently provides about this option:

-**Xmx***n*

Specify the maximum size, in bytes, of the memory allocation pool. This value must be a multiple of 1024 greater than 2MB. Append the letter k or K to indicate kilobytes or the letter m or M to indicate megabytes. The default value is 64MB. Examples:

```
-Xmx83886080
-Xmx81920k
-Xmx80m
```

The term "viewmgr" is used throughout this chapter to represent the Java command for View Manager.

The syntax for the View Manager command line uses the following conventions:

Table 8.4 Syntax for the Command Line

Convention	Description
Bold	Information that must be used exactly as shown.
Italics	Information that will be replaced with the names of your files, child folders, etc.
[ ]	Square brackets surround optional syntax.
|	A vertical bar separates mutually exclusive choices. You select only one of the choices.
//	Because the View Manager command line performs many operations and has many options, comments are intermixed with the syntax. Comments start with two forward slashes (//).
	Each command you create should be one line, despite the fact the syntax is broken up into separate lines in the following section.

View Manager Command Line Syntax

The following section presents the View Manager command line syntax. Although the syntax normally appears on just one line (to be sent to the operating system), it is presented here in pieces, for better understanding.

```
//specifying the command
//replace viewmgr with whatever is appropriate for your platform
   viewmgr

//specifying the server
-s "userName:[password]@hostName:endpoint"

//server properties (password, compression, and encryption)
//-pwdfile is shown here because it may be used with options that specify
//the server. -pwdfile is also shown with the -open option with which it is
//required.
   [-pwdfile "filePath"] [-cmp]
   [-encrypt encryptionType]

//specifying the projects
//-p is an alternative when the source and target project have the same
//name
   [-srcproject "projectName"
   -tgtproject "projectName" |
   -p "projectName"]

//specifying the views
//-v is an alternative when the source and target view have the same
//path
   [-srcview "viewPath" -tgtview "viewPath" |
   -v "viewPath"]

//specifying the folders
//-f is an alternative when the folders have the same path
   [-srcfolder "folderPath" -tgtfolder "folderPath" |
   -f "folderPath"]

//specifying view configuration when view is source object
   [-srclabel "viewLabel" | -srcstate "promotionState"]

//including child views/folders when view/folder is source object
   [-cf | -cv]

//specifying run operation
//when -bi, -nomerge, and -copy are not specified, View Manager
```

```
//performs a source-to-target synchronization by default
//-bi is for bidirectional synchronizations
//-nomerge is for missing-objects-only synchronizations
//-copy is for copy-to-target operations
   [-bi | -nomerge | -copy] [-preview]

//controlling when and if the merge utility is used for text files;
//controlling how binary files are “merged”
   [-batch | -postpone] [-noautomerge]
   [-usesource [text | binary] |
   -usetarget [text | binary]]

//merge utility options used with preview operations only
   [-b | -diffutil "utilityPath" -diffopts "utilityOptions"]

//merge utility options used with source-to-target or bidirectional operations only
   [-b | -mergeutil "utilityPath" -mergeopts "utilityOptions"]

//merge options when recording and tracking merges
   [-recordmp] [-usemp| -usemptgt]

//breaking locks so a merged file can be checked in
   [-breaklock]

//using the properties files
   [-open "filePath" -pwdfile "filePath" [-resume] | -save "filePath"]

//replicating times, users, view labels and promotion states, revision label, history of
files
   [-reptimes [-vlabels [-pstates]]] [-repusers] [-rlabels] [-hist]

//managing locks

//locking/unlocking the server (notice the "s" at the end, despite the fact
//you are locking/unlocking only one server)

[-lockservers [-unlockservers]]

//leaving text or binary files locked after the run

[-lock [text | binary]]

//missing objects options (applies to source-to-target, bidirectional, and missing-
objects-only run operations)
//You might use each command twice: once for source and once with target

[-sp [source | target]] [-sv [source | target]]

[-sf [source | target]] [-si [source | target]]

[-df [source | target]] [-di [source | target]]

//generic options

[-nologo] [-?]
```

Option	Description
`-?`	At the command line, lists the command syntax and describes each option.

`-b`	When comparing two lines of text files, ignores trailing white space and treats all other strings of white space as equal in length. For example, the following lines are equivalent:
	`" hi mom  "`
	`"  hi  mom"`
	This option is meaningless if you use -nomerge.

`-batch`	Enables you to run View Manager uninterrupted and worry about files that have conflicts or differences later. No manual merging is done. If you use -noautomerge with -batch, no files are merged, manually or automatically. View Manager sends error messages to standard error for unresolved text files. (This action is in addition to the error messages sent by View Manager for binary files.) This option is meaningless if you use -nomerge. Without this option or the -postpone option, the merge utility appears whenever View Manager encounters text files that require manual merging.
`-bi`	Specifies a bidirectional synchronization.
`-breaklock`	Enables View Manager to break another user's lock on a file if the file needs to be checked in. You must have the correct access rights to do this.
`-cf`	When a folder is selected as the source object, this option includes the child folders of that folder in the run operation. Without this option, only the specified folder and its file contents are used in the run operation.
`-cmp`	Compresses all data sent between the workstation and the server and decompresses it when it arrives. Without this option, the default is no compression. Compression is most useful and appropriate when the client and server communicate over a slow connection. To determine whether to use compression, you must consider whether the time spent compressing and decompressing data is better than the longer transfer time of uncompressed data sent over the slow connection. A small test case may be helpful. The server may enforce a higher level of compression than you specify.
`-copy`	Indicates that the run operation is a copy-to-target operation.
-cv	When a view is selected as the source object, includes its child views in the run operation
`-df`	Deletes folders from the specified object if they do not exist in the unspecified object. The full syntax is: `-df source | target` Use -df to indicate which object will change as a result of using this option. For example, `-df target` deletes folders from the target object that do not exist in the source object. Deleting a folder deletes its entire contents: child folders and files, recursively. You may use this option more than once in a command if you want to change both the source and the target. No folder is deleted based on this option unless its parent (or an equivalent object) exists in both the source and target objects. If a folder is the selected source or target object, that folder is *not* missing. If a folder is the root folder of the selected view or the root folder of the root view of the selected project, it is also *not* missing. This option has no effect when used with -copy. Use it with synchronizations: missing-objects-only (-nomerge), bidirectional (-bi), or source-to-target (the default) operations.

Option	Description
`-di`	Deletes files from the specified object if they do not exist in the unspecified object. This option applies to files that are in folders that already exist in both objects. The full syntax is: `-di source | target` Use -di to indicate which object will change as a result of using this option. For example, `-di target` deletes files from the target object that do not exist in the source object. You may use this option more than once in a command if you want to change both the source and the target. No file is deleted based on this option unless its folder (or an equivalent folder) exists in both the source and target objects. This option has no effect when used with -copy. Use it with synchronizations: missing-objects-only (-nomerge), bidirectional (-bi), or source-to-target (the default) operations.
`-diffopts`	Must be used with the -diffutil option to provide the command-line options that go with that utility. The full syntax is: `-diffopts "`*utilityOptions*`"` You can include the following terms in *utilityOptions*: ■ $file1, as a place holder for the path to the source file to be compared. ■ $file2, as a place holder for the path to the target file to be compared. For example, suppose the comparison utility's path is "D:\Programs\Comparison Utility 8\Diffutil.exe" and the options to be used are "$file1 $file2". View Manager assumes that the options are appropriate for the utility to be used. Each time that View Manager needs to call the comparison utility, it replaces the terms that start with $ with the actual paths to those files and asks the operating system to execute the completed command. On UNIX systems, you need to use the escape character "\" in front of "$". For example, you would use "\$file1" instead of "$file1".
`-diffutil`	Specifies the use of a comparison utility other than Visual Diff by providing the complete path to that utility. The full syntax is: `-diffutil "utilityPath"` This option must be used with -diffopts to provide the command-line options for the comparison utility. It is ignored unless the operation is a source-to-target or bidirectional preview operation.

Option	Description
`-encrypt`	Encrypts all data sent between the workstation and the server and unencrypts it when it arrives. Without this option, no encryption occurs. Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines. The full syntax is: `-encrypt` *encryptionType* The types of encryption are: **RC4** RSA R4 stream cipher (fast) **RC2_ECB** RSA R2 block cipher (Electronic Codebook) **RC2_CBC** RSA R2 block cipher (Cipher Block Chaining) **RC2_CFB** (Windows only) RSA R2 block cipher (Cipher Feedback) The encryption types are ordered from fastest to slowest. Each of the slower encryption types is safer than the one preceding it. For platforms other than Windows, the public and private keys used in the encryption process are not created automatically. They are stored in an options file named .starteam that is usually stored in your home directory. It contains a variable or shell variable called keyfile. This variable specifies the location of the file containing the public and private keys. If you do not specify the keyfile variable, an error occurs. When you specify the keyfile variable, but the file does not exist, the client generates a random pair of keys, creates the file, and stores the keys in it. Be sure to secure the .starteam file. For example, in UNIX, only its owner should be able to read it.
`-f`	When both the source and target folders have the same path, you can use -f instead of -srcfolder and -tgtfolder. The full syntax is: `-f` "*folderPath*" To designate the root folder use: `"/"` The folder hierarchy should never include the root folder. For example, if the root folder of the view is StarDraw and the hierarchy to your folder is "StarDraw/SourceCode/Client", use only "SourceCode/Client". For example: `-f "Source Code/Client"`
`-hist`	Copies the history of a file along with the file. If you are replicating a parent view and at least one of its child views, you must copy the history. If you do not include -hist, View Manager notifies you of this fact and asks if you want to proceed with -hist or cancel the operation.
`-lock`	Enables you to exclusively lock unresolved text and/or binary files. The full syntax is: `-lock [text | binary]` Use text to indicate that only unresolved text files are to be locked. Use binary to indicate that only unresolved binary files are to be locked. Use neither to lock all unresolved files.
`-lockservers`	Locks the server. Notice that this option ends in "s" even though you have only one server. This is for backwards compatibility with another set of commands.

Option	Description
`-mergeopts`	Must be used with the -mergeutil option to provide the command-line options that accompany that utility. The full syntax is: `-mergeopts "utilityOptions"` You can include the following terms in *utilityOptions*: $branchtip, as a place holder for the path to the tip revision of the source file to be merged. $usertip, as a place holder for the path to the target file to be merged. $basefile, as a place holder for the path to the common ancestor for the $branchtip and $usertip files. $resultfile, as a place holder for the path to the file that will store the output from the merged file. For example, suppose the merge utility's path is "D:\Programs\Merge Utility 5\Mergeutil.exe" and the options to be use are "-s $branchtip $usertip $basefile $resultfile". View Manager assumes that the options are appropriate for the utility to be used. Each time that View Manager needs to call the merge utility, it replaces the terms that start with $ with actual paths to the files and asks the operating system to execute the completed command. On UNIX systems, use the escape character "\" in front of "$". For example, you would use "\$branchtip" instead of "$branchtip".
`-mergeutil`	Specifies the use of a merge utility other than Visual Merge by providing the complete path to that utility. The full syntax is: `-mergeutil "utilityPath"` This option must be used with -mergeopts to provide the command-line options for the merge utility. It is ignored unless the operation is a source-to-target or bidirectional synchronization. It does not apply to previews.
`-noautomerge`	Stops View Manager from merging text files automatically when they have no conflicts. It displays the merge utility unless -usesource and -usetarget make that unnecessary. When this option is not used, by default, View Manager merges text files that have no conflicts.
`-nologo`	Suppresses the copyright notice.
`-nomerge`	Does not merge files; therefore this command is equivalent to performing a Missing Objects Only run operation. Use this option when you want to add/share or delete items without merging files. (Views are added; folders and files are shared.) Without this option, text files in the source view are merged with the corresponding files in the target view, if a merge is required. If you use -nomerge, all other merge options, such as -batch and -recordmp, are ignored.
`-open`	Specifies the path to the properties file (.srm) that contains the options to be used with this command. The full syntax is: **-open** "*filePath*" If you use -open, you must use -pwdfile because the password is not stored in the properties file. You must also specify -bi, -nomerge, or -copy unless you want the default source-to-target synchronization.

`-p`	When both the source and target projects are the same, you can use -p instead of -srcproject and -tgtproject.
	The full syntax is:
	`-p "projectName"`
	`For example:`
	`p "Big Project"`

Option	Description
`-preview`	Performs a preview operation. No items are merged, although they can be compared. Data about the operation appears in the automatically generated .log files. If you use -preview with -recordmp, -batch, -usesource, or -usetarget, -preview takes precedence.
`-postpone`	Causes manual merging to occur at the end of the run. When used in combination with -noautomerge, it postpones processing of all text files that have differences. Otherwise it postpones the processing of all text files that have conflicts. Without this option or the -batch option, the merge utility appears whenever View Manager encounters text files that require manual merging.
`-pstates`	Copies all promotion state definitions from one server to the other. Depending on your missing-object option settings and the type of operation, promotion state definitions can be copied from the source to the target and the target to the source. Promotion states are valid at the view level. If, for example, you select a folder as the source object, the promotion state information from the view containing that folder is copied to the target view in which the folder is replicated. If you select a project as a source object, the promotion state information from each view in that project is replicated to the corresponding view in the replicated project. If two corresponding views have promotion states with the same name, no copying occurs. If the source and target objects are the same view or folders in the same view, this object is not necessary. This option must be used with -vlabels or it is ignored.

`-pwdfile`	Specifies the complete path to a file that stores your password. This option provides the password normally supplied as part of the -s option. It prevents others from seeing your password on the command line. Leading and trailing white space is ignored.
	The full syntax is:
	`-pwdfile "`*fileName*`"`
	When you use a password file, the syntax of -s changes to remove the password but not the colon (:) between your name and the password:
	`-s "`*userName*`:@`*hostName*`:`*endpoint*`"`
	With -open, this option is required.

`-recordmp`	Records a merge point is recorded for each pair of files (in a given direction) if any of the following is true: ■ A merged file is created from the pair of files and is checked in. ■ The file generated by the merge utility does not need to be checked in. For example, this is the case in a source-to-target operation if the contents of the merged file match those of the target file. ■ No merge occurs because you used the -usemp or -usemptgt option and the ancestor file has not changed since the last merge in this direction. For example, suppose you perform a source-to-target operation using -usemp. The -usemp option uses the source file from the last recorded merge as the ancestor file during the merge. If the source file has not changed since the last merge, no merge occurs because nothing needs to be added to the target file. ■ The -usesource or -usetarget option has been used with text files, which makes a merge unnecessary. This option is meaningless if you use -nomerge or -preview.
`-reptimes`	Replicates the creation time property of objects as well as the objects. Without this option, the new object has the current time.

Option	Description
`-repusers`	Copies all users from source server to target server. If objects are copied from the target to the source, it also copies all users from the target to the source server. Newly created users have suspended accounts with the password “password,” regardless of their passwords on the original server. If View Manager needs to log on as an existing user, that user’s password is changed to “password.” Enables View Manager to include user names as the “Created By” property of created objects. View Manager must log on as a user to create an object with that user’s name.
`-resume`	In conjunction with the -open option, enables you to resume a run operation that halted. For example, the operation may have ended because you cancelled it, the number of retries was exhausted, or a fatal exception occurred. If you list all the options at the command line (instead of using -open to specify an operation’s options), it is a good idea to include -save in case you need to resume.
`-retries`	Number of times View Manager will attempt to resume an operation after becoming disconnected from the server. The range is 0 to 99. The default is 3. Use this command when you want more or fewer retries. The difference between -retries and -resume is that -retries continues an operation that halts during the execution of a command and -resume continues an operation that halted during the execution of a previous command. If you set -retries to 3, View Manager will reconnect three times during the command. If the command halts again, no additional retries are attempted. However, you can use the -open and -resume options in a new command that starts where the first command ended.
`-retrywait`	Number of minutes to wait before resuming an operation after becoming disconnected from the server. The range is 0 to 99. The default is one minute, but it may take a few minutes for StarTeam Runtime to detect the disconnect. Use this command when you want more or fewer minutes between retries.
`-rlabels`	Copies the revision label information associated with a file revision as well as the file revision. Although a file revision can have any number of revision labels, a revision label can be assigned to only one file revision in a specific file’s history at a given time. When a file revision is replicated, whether from source-to-target or vice versa, the immediate transaction takes precedence over past transactions. For example, suppose that the revision label “xxx” is attached to revision 4 and that revision 4 is being copied from the source to the target. If the target’s revision 2 already has the same revision label, the label will be moved from revision 2 to revision 4.
`-s`	Specifies your name and password and the server’s address and endpoint. The full syntax is: `-s "`*userName*`:`*password*`@`*hostName*`:`*endpoint*`"` For example: -s "JMarsh:password@orion:49201" If the user name is omitted, the current user’s logon name is used. The user name in the example is “JMarsh” If the password is omitted, you are prompted to enter the password. The password in the example is “password” If the host name is omitted, the default is localhost. The host name in the example is “orion”. The default endpoint (port) for TCP/IP is 49201.The endpoint in the example is 49201.

Option	Description
`-save`	Specifies the path to the file in which the options used with this command are saved. Unlike the properties files saved from View Manager graphical user interface, only one set of command-line options can be saved. The full syntax is: **-save** *"filePath"*
`-sf`	Shares folders to a specified object from an unspecified object if they do not already exist in the specified object. The full syntax is: `-sf source | target` You indicate which object will change as a result of using this option. For example, `-sf target` shares folders to the target object that exist only in the source object. The command includes the folders, its child folders, and files, recursively. You may use this option more than once in a command if you want to change both the source and the target. No folder is shared based on this option unless its parent (or an equivalent object) exists in both the source and target objects. If a folder is the selected source or target object, that folder is *not* missing. If a folder is the root folder of the selected view or the root folder of the root view of the selected project, it is also *not* missing. This option has no effect when used with -copy. Use it with synchronizations: missing-objects-only (-nomerge), bidirectional (-bi), or source-to-target (the default) operations.[1]
`-si`	Shares items to a specified object from an unspecified object if they do not already exist in the specified object. The full syntax is: `-si source | target` You indicate which object will change as a result of using this option. For example, `-si target` shares files to the target object that exist only in the source object. You may use this option more than once in a command if you want to change both the source and the target. No file is shared based on this option unless its folder (or an equivalent folder) exists in both the source and target objects. This option has no effect when used with -copy. Use it with synchronizations: missing-objects-only (-nomerge), bidirectional (-bi), or source-to-target (the default) operations.
`-src folder`	Specifies the source folder. The full syntax is: `-srcfolder "`*folderPath*`"` To designate the root folder use: `"/"` To designate a folder other than the root folder, omit the root folder from the folder hierarchy. For example, if the root folder of the view is StarDraw and the hierarchy to your folder is "StarDraw/SourceCode/Client." use only "SourceCode/Client". For example: -srcfolder "Source Code/Client"
`-srclabel`	If the source object is a view, you can configure (or roll back) that view to a specified view label and replicate the view as it was at that point in time. The full syntax is: `-srclabel "`*viewLabel*`"` This option is mutually exclusive with -srcstate. When neither option is used, the view configuration is the current configuration.

Option	Description
`-srcproject`	Specifies the source project. The full syntax is: `-srcproject "`*projectName*`"` For example: -srcproject `"Big Project"`
`-srcstate`	If the source object is a view, you can configure (or roll back) that view to the view label assigned to a specified promotion state and replicate the view as it was at that point in time. The full syntax is: `-srcstate "`*promotionState*`"` This option is mutually exclusive with -srclabel. When neither option is used, the view configuration is the current configuration.

`-srcview`	Specifies the source view or the path to that view from the root view. When specifying the path, the root view is always the first view in the path.
	The full syntax is:
	`-srcview "`*viewPath*`"`
	For example:
	`-srcview "Big Project/Child View"`
	The root view in the example is Big Project, while Child View is a child of the root view.
	If two child views have the same name, *viewPath* should contain the path from the root view to the child view, to eliminate ambiguity. Otherwise, the command uses the first view it locates that has the correct name.

`-sv`	Adds views to a specified object from an unspecified object if they do not already exist in the specified object. The full syntax is: `-sv source | target` You indicate whether the source or target object will change as a result of using this option. For example, `-sv target` adds views to the target object when they exist only in the source object. This option includes the view, its child views, and their folders and files, recursively. You may use the option more than once in a command if you want to change both the source and the target. No view is added based on this option unless its parent (or an equivalent object) exists in both the source and target objects. If a view is the selected source or target object, that view is *not* missing. If a view is the root view of the selected project, it is also *not* missing. This option has no effect when used with -copy. Use it with synchronizations: missing-objects-only (-nomerge), bidirectional (-bi), or with source-to-target (the default run operation). No similar deletion option exists. Views cannot be deleted using View Manager. Use StarTeam for this purpose.

`-tgtfolder`	Specifies the target folder.
	The full syntax is:
	`-tgtfolder "`*folderPath*`"`
	To designate the root folder use: `"/"`
	To designate a folder other than the root folder, omit the root folder from the folder hierarchy. For example, if the root folder of the view is StarDraw and the hierarchy to your folder is "StarDraw/SourceCode/Client." use only "SourceCode/Client".
	For example:
	`-tgtfolder "Source Code/Client"`

Option	Description

`-tgtproject`	Specifies the target project.
	The full syntax is:
	`-tgtproject "`*projectName*`"`
	For example:
	`-tgtproject "Big Project"`

`-tgtview`	Specifies the target view or the path to it from the root view. When specifying the path, the root view is always the first view in the path.
	The full syntax is:
	`-tgtview "`*viewPath*`"`
	For example:
	`-tgtview "Big Project/Child View"`
	The root view in the example is Big Project., while Child View is a child of the root view.
	If two child views have the same name, *viewPath* should contain the path from the root view to the child view, to eliminate ambiguity. Otherwise, the command uses the first view it locates that has the correct name.

`-unlockservers`	Unlocks the server after the operation (and manual merges) are completed. Notice that this option ends in "s" even though you have only one server. This is for backwards compatibility with another set of commands.
`-usemp`	Bases the merge on the source file from the last recorded merge, as well as on the pair of files being merged (mp stands for merge point). Without -usemp or -usemptgt, the most recent common ancestor is used to perform the three-way merge. The common ancestor is a historical revision with the same revision number in both views. This option is meaningless if you use -nomerge.
`-usemptgt`	Bases the merge on the results of the last recorded merge, as well as on the pair of files being merged. The results from the last recorded merge were checked in to the target view. (mp stands for merge point and tgt stands for target.) Without -usemp or -usemptgt, the most recent common ancestor is used to perform the three-way merge. The common ancestor is a historical revision with the same revision number in both views. This option is meaningless if you use -nomerge.
`-usesource`	Substitutes the source file for the merged file when a manual merge would have occurred. The merge utility does not display. When this option is used for text files, if -noautomerge is included, the substitution is for text files that have differences. Otherwise, the substitution is for text files that have conflicts. The full syntax is; `-usesource [text | binary]` Use text to indicate text files. Use binary to indicate binary files. Use neither to indicate both text and binary files. This option is meaningless if you use -nomerge. View Manager records the merge point if -recordmp is used and moves on to the next pair of files.

Option	Description
-usetarget	Substitutes the target file for the merged file. The merge utility does not display. When this option is used for text files, if -noautomerge is included, the substitution is for text files that have differences. Otherwise, the substitution is for text files that have conflicts. The full syntax is; **-usetarget** [text | binary] Use text to indicate text files. Use binary to indicate binary files. Use neither to indicate both text and binary files. This option is meaningless if you use -nomerge. View Manager records the merge point if -recordmp is used and moves on to the next pair of files.
-v	Used instead of -srcview and -tgtview when both the source and target view have the same path. The full syntax is: **-v** "*viewPath*" For example: -v "Big Project/Child View" The root view in the example is Big Project, while Child View is a child of the root view. *viewPath* can be the name of the view or the path to it from the root view. If two child views have the same name, *viewPath* should contain the path from the root view to the child view, to eliminate ambiguity. Otherwise, View Manager uses the first view it locates that has the correct name.
-vlabels	Copies all view labels from the source object to the target object. If objects are copied from the target to the source, it also copies all view labels from the target to the source server. View labels are valid at the view level. For example, if you select a folder as the source object, the view label information from the view containing that folder is copied to the target view in which the folder is replicated. If you select a project as a source object, the view label information from each view in that project is replicated to the corresponding view in the replicated project. If two corresponding views have view labels with the same name, no copying occurs. Although a folder or file can have any number of view labels assigned to the same revision, a given view label can be assigned to only one file revision in a file's history at a given time. When a file revision is replicated, whether from source-to-target or vice versa, the immediate transaction takes precedence over past transactions. For example, suppose that the view label "build xxx" is attached to revision 4 and that revision 4 is being copied from the source to the target. If the target's revision 2 already has the same view label, the label will be moved from revision 2 to revision 4.

1. Be aware of the effects of using -sf and -df or -si or -di together. For example, if you use -sf source with -df source, -sf source will share folders to the source that were only in the target, while -df source will remove folders from the source that were not in the target. These two options work with different sets of folders. Using -sf source and -df target together moves folders to the source that were not in the source but were in the target. Then those folders are removed from the target.

Logging Screen Output

During a View Manager command, data appears on the screen explaining what is happening. This data has been sent to standard output or to standard error. You can have this information logged instead of displayed.

On Windows, if you end the command with "1>2& viewmgr.log", the data that is sent to either standard output and standard error appears in viewmgr.log. Ending the command with ">viewmgr.log" logs only the data from standard output.

The screen output is not critical data, just a list of the objects that have been processed. The automatic log files are much more comprehensive.

Learning from Examples

The commands in this section are representative of those you would use when synchronizing files.

Example 1

The following View Manager command merges files in a source view with files in the target view. The source view is the root view of the project named Big Project. The target view is a child of that root view named Child View.

Because the command includes the -usemp option, the command determines whether two corresponding files need to be merged based on whether the item in the source view (in this case, Child View) has changed since the last recorded merge.

Because the command includes the -recordmp option, the merge will be recorded as a merge point, and the merged file will be used the next time the same pair of files is merged using this command, the same source and target views, and the -usemp option.

In addition to the merge operations, this command adds folders and files from the source view to the target view. The added folders and files are those that exist in the source view but are missing from the target view.

```
viewmgr -s "JMarsh:password@orion:49201" -p "Big Project" -srcview "Big
Project" -tgtview "Big Project/Child View" -recordmp -b -usemp -si source -sf
source
```

Example 2

The next command merges files using the same two views. However, because neither the -usemp nor -usemptgt option is used, recorded merge points are ignored. Because the -recordmp option is not used, this merge does not record a merge point for future merges either.

When the source file and the target file are different, they are merged automatically if they have no conflicts. This command can be run overnight with no supervision because the -batch option bypasses files that have conflicts.

```
viewmgr -s "JMarsh:password@orion:49201" -p "Big Project" -srcview "Big
Project" -tgtview "Big Project/Child View" -batch -b
```

Without -batch, View Manager displays the Visual Merge application for files that have conflicts so that you can resolve the conflicts.

If the person running the command is not the person who can resolve the conflicts, using -batch enables others to merge the files with conflicts later by using the View Manager's graphical user interface or command line.

Example 3

The next command is similar to the previous one. However, instead of ignoring files with conflicts until later, this command checks the source file into the target view as though it were the merged file.

```
viewmgr -s "JMarsh:password@orion:49201" -p "Big Project" -srcview "Big
Project" -tgtview "Big Project/Child View" -recordmp -usesource
```

Chapter

9

Using Views

When you create a project, an initial or root view of that project is also created. This initial view, which has the same name as the project, consists of the root folder and a hierarchy of child folder. It is always read/write.

To accommodate both user and project needs, however, the application enables you to create additional views of a project based on this view. These additional views may contain some or all of the contents of the original view and may behave differently.

For example, you might use views to:

- Implement the same folder hierarchy for both the 2.5 release and the 3.0 release of a product. The easiest way to do this is by creating a new view for the 3.0 release based on the view used for the 2.5 release.
- Limit the portion of the project that certain team members see. Developers might need to see only the project's source code folder and its child folders; marketing personnel might need to see only the marketing folder and its child folders; and so on. Each of these views can have a different folder as its root.
- Support branching and parallel development. By branching files and other data in a new view, you organization can start to work on the 2.0 version of a product without hampering the creation of service packs for the 1.0 version.

Views typically have names such as Baseline, 5 Maintenance, Special 5 for Australia, and 5.0 New Development. They represent configurations of items and support different development baselines of the same code base. If desired, views can be compared and merged. For example, you might want eventually to merge files from 5 Maintenance and 5.0 New Development into the Baseline view.

This chapter contains the information about views that administrators need to know but that most users probably do not require. More information about views can be found in “Using Projects and Views” in the *StarTeam User’s Guide*.

Understanding Different Types of Views

The root view displays when you highlight View > Select Configuration > Current Configuration on the menu. It is considered a dynamic view, because it shows all items in the project as they change, making it ideal for collaborative development. Once it exists, however, you can use the New View Wizard to create other views based on it .

These additional views are created with the StarTeam New View Wizard. The options selected when you use the New View Wizard determine their characteristics. For example, these views can be read-only or read-write, be frozen or not, and branch or not.

The remainder of this section reviews additional types of views and their uses. For a complete description of each type of view and its features, see Table 9.1, “Creating Views with Desired Characteristics,” on page 141. This table also lists the settings that must be selected in the New View Wizard to create the views.

Branching Views (Branch All, Branch None)

Branching views must be created from an existing view that permits branching (i.e., not a reference view). In special circumstances, you may even create a branching view that contains none of the contents of its parent. See “Creating Additional Views” on page 144.

In a branching view, not all folders or items in a folder need to branch. Also, no branching occurs unless an item that changes has its branching behavior set to branch on change.

Important Remember that only folders, files, and change requests can branch. Requirements, tasks and topics never branch. Changes to such items can be made in the view in which they were created, in a view in which they have been moved or shared, or in floating child views that are not read-only. New revisions of such items can be reviewed from any floating child view.

After an item branches in the child view, its revision number indicates a new branch. For example, the item in the parent view might have the revision number 1.7, while the revision number for the corresponding item in the child view becomes 1.7.1.0. The only way to make items identical after they have branched is to merge them. For more information about merging items in views, see Chapter 10, “Comparing and Merging Views,” on page 149

For reason’s of safety, deletions made in the parent view are not propagated to the child view and vice versa. If you want to delete a folder or item from all related views, you must delete it in each of the views.

Similarly, moving a folder or item is considered a copy operation followed by a delete operation. Because the delete operation is not propagated, the view in which the move was made has one copy of the folder or item in the new location, while related views have two copies of the folder or item, one in the original location and one in the new location. Thus, a move is equivalent to a share in the related views.

A branching view should use a different working folder than that of its parent view. Using the same working folder for both views is not only confusing but can create serious problems:

- Changes to files in one view may be overwritten when you check out files from the other view. This problem occurs because the two files have the same name and use the same working folder.
- File status can be often wrong or, at least, misleading. For example, the application may report a the status of a file in one view as Unknown, Out of Date, or Merge, when it is really Missing and the file from the other view with the same name is in the working folder.

Possible Uses of Branching Views

With this application, your company can use labels, promotion states, and branching views for continuous baseline development during release engineering activities. The following steps outline this process:

1. Create a view label when development is complete. For example, you might name the view label Release Candidate.
2. Select View > Promotion from the menu bar.
3. Use the *Promotion* dialog to create various promotion states (such as Testing, QA, and Release.)
4. As work progresses, use the *Promotion* dialog to promote the Release Candidate view label through the various states, until it reaches the final Release state.
5. Create a branched view derived from the current view, and assign this configuration to a promotion state. In this case, the desired promotion state is Release.

   A separate baseline is created for release integration and engineering, while development continues in the main baseline.

This application also enables users to create a separate baseline for independent development efforts. This baseline can be based on the main baseline in its current or an earlier configuration.

When developing an independent product that is based on an existing product, Borland recommends creating a branching view. Each item in the new view should branch when that item changes in the new view.

The following steps outline this process:

1. From the main baseline, create a new view from the current view.
2. Permit items to branch, and branch all items.
3. Using the *New View Wizard,* create the view. Give the new view a name and a new working folder. Select a configuration based on the specific date and time, label, or promotion state.
4. Check out the files in the new view, change them as appropriate, and check them back in.

Tip When you create a branching view based on a specific label, promotion state, or time, each item in the child view has its configuration set to the corresponding time. If changes are made to the parent item before the child item branches, you can advance the configuration time to include the changes in the parent.

For example, if an item is configured to June 5, 2003, 4:44 P.M. and changes are made to the parent item on June 8th and July 14th, you can advance the configuration time of the child item to an appropriate time on July 14th, thus including the changes in the child view.

Folders and Items in Both the Parent and the Branching View

Depending on the folders included in the new branching view, certain items from the parent view appear in the new branching view. However, these “inherited” folders and items will have no label, as neither view labels nor revision labels move from view to view. Also, if you roll back a view to an earlier point in time, you will no longer see the folders and/or items that have been moved.

As a result, the workflow for change requests is affected in the following ways:

- If the Last Build Tested and the Addressed In Build fields in a change request have no values when the change request branches, their workflow is specific to the new view only.
- If the Last Build Tested and the Addressed In Build fields in a change request have build labels as their values (i.e., these fields are not empty or do not contain the value “Next Build”) when the change request branches, the branched change request retains those values. In the new view, these values can be changed, but only to the names of build labels that exist in the new view.

- If the Addressed In Build field contains the value “Next Build” when the change request branches, the “Next Build” value is replaced by the name of the next build label created in the parent view, not the next build label created in the new view.

Reference Views

A read/write reference view is usually created to show only a subset of the folders found in the parent view. For example, suppose a project has StarDraw as its root folder and contains several child folders, including one named SourceCode. You can create a read/write reference view for developers that has SourceCode as its root folder, allowing them to work with only the parts of the project that are relevant.

Read-Only Reference Views

Read-only reference views are reference views to which users cannot make changes. That is, users can check out existing files, but they cannot edit them or check them back in.

The configuration of read-only reference views can be either:

- Floating, which stays current with the parent view. For example, a floating view could be created for team members who are monitoring changes to a project. Read-only reference views that float can be rolled back.
- Frozen (or fixed) at a particular date and time in the parent view. For example, a frozen view might be created as a milestone against which to change further progress on a project. Read-only reference views that are frozen cannot be rolled back.

The access rights of a read-only reference view are the same as those used in its parent view. If you change the access rights in the parent, the rights are automatically changed in the read-only view, even if that view is configured to a point in the past.

Non-Derived Views (Blank Branching Views)

Non-Derived views create a branch but start with no files and no child folders. Essentially, they are not derived from an existing view, even though a parent is selected for them and they display in the view hierarchy. Folders and items are usually added to, moved to, or shared with these views.

Table 9.1 Creating Views with Desired Characteristics

Desired Characteristics	Options to Set
Branching: Branch All (Floating) **New Items:** New items in the child view do *not* appear in the parent view; new items in the parent view appear in both views if they are in the subset accessed by the child view. In the child view, new items from the parent have the Branch on Change check box selected. **Existing Items:** Changes to existing items in the parent view appear in the child view until the corresponding item in the child view branches. Changes to existing items in the child view can appear in the parent view, but only if the Branch on Change check box for that item is cleared. However, items deleted from one view are not deleted from the other. **Item Behavior:** Branch on Change check box is enabled and selected for all items that can branch. Any change to a child item that can branch results in the branching of that item, unless its Branch on Change check box has been cleared. After branching, the box is disabled.	**View Type:** Branch All **Root Folder:** Selected from parent view **Working Folder:** Should be different from that of the parent, to avoid conflicts **Configuration:** Floating **Note:** If users are likely to perform many move and share operations, using branching, floating views can result in multiple unwanted references to the same folders or items, causing confusion. Also, if a change request has “Next Build” in the Addressed In Build field when the child view is created, “Next Build” will be replaced by the parent’s next build label, unless the change request is first branched.
Branching: Branch All (Not Floating) **New Items:** New items in the child view do *not* appear in the parent view; new items in the parent view do *not* appear in the child view. **Existing Items:** Existing items in the child view are the same as in parent view at the time of configuration, until the item in the child view branches. Changes cannot be made to an item in the child view if the change does not result in branching. **Item Behavior:** Branch on Change check box is enabled and selected for all items that can branch. Any change to a child item that can branch results in the branching of that item, unless its Branch on Change check box has been cleared. After branching, the check box is disabled.	**View Type** Branch All **Root Folder:** Selected from parent view **Working Folder:** Should be different from that of the parent, to avoid conflicts **Configuration:** Other than floating (Labeled, Promotion state, or As of specific date) **Note:** If a change request has “Next Build” in the Addressed In Build field when the child view is created, “Next Build” will not be replaced by a build label until the change request branches.

Table 9.1 Creating Views with Desired Characteristics (continued)

Desired Characteristics	Options to Set
Branching: Branch None (Floating) **New Items:** New items in the child view appear in the parent view; new items in the parent view appear in the child view if they are in the subset accessed by the child view. In the child view, new items from the parent have the Branch on Change check box cleared. **Existing Items:** Existing items are the same in the child view as in the parent view; they can be changed in either the parent or child view until the item in the child view branches. However, items deleted from one view are not deleted from the other. **Item Behavior:** Branch on Change check box enabled, but cleared. No change to a child item that can branch results in branching until the Branch on Change check box is selected. After branching, the check box is disabled.	**View Type:** Branch None **Root Folder:** Selected from parent view **Working Folder:** Should be different from that of the parent, to avoid conflicts **Configuration:** Floating **Note:** If users perform many move and share operations, using branching, floating views can result in multiple unwanted references to the same folders or items, causing confusion. Also, if a change request has "Next Build" in the Addressed In Build field when the child view is created, "Next Build" will be replaced by the parent's next build label, unless the change request branches in the child view.
Branching: Branch None (Not Floating) **New Items:** New items in the child view do *not* appear in the parent view; new items in the parent view do *not* appear in the child view **Existing Items:** Existing items are exactly the same in the child view as in the parent view at the time of configuration, until the item in the child view branches. Changes cannot be made to an item in the child view if that change does not result in branching. **Item Behavior:** Branch on Change check box enabled, but cleared. No change to a child item that can branch results in branching until the Branch on Change check box is selected. After branching, the box is disabled.	**View Type:** Branch None **Root Folder:** Selected from parent view **Working Folder:** Should be different from that of the parent, to avoid conflicts **Configuration:** Other than floating (Labeled, Promotion State, or As of specific date) **Note:** If a change request has "Next Build" in the Addressed In Build field when the child view is created, "Next Build" will not be replaced by a build label until the change request branches.
Reference view Also called Read/Write Reference view **New Items:** New items in the child view appear in both views; new items in the parent view appear in both views if they are in the subset accessed by the child view. **Existing Items:** Existing items are the same in the child view and the parent view; they can be changed from either view. **Item Behavior:** Branch on Change disabled **Note:** Labels created and objects deleted in the child view appear and disappear in the parent view; this is not true for other types of child views.	**View Type:** Reference **Root Folder:** Selected from parent view **Working Folder:** Usually the same as that of the parent view **Configuration:** N/A; always floats.

Table 9.1 Creating Views with Desired Characteristics (continued)

Desired Characteristics	Options to Set
Read-only reference view (Floating) **New Items:** New items cannot be added to child view; new items in parent view appear in child view. **Existing Items:** Existing items are the same in the child view and the parent view; they can be changed only from the parent view. **Item Behavior:** Branch on Change disabled.	**View type:** Read-only Reference **Root Folder:** Selected from parent view **Working Folder:** Usually the same as that of the parent view **Configuration:** Floating. These views can be rolled back.
Read-only reference view (Frozen) Also called a snapshot view **New Items:** New items cannot be added to the child view; new items in parent do not appear in the child view. **Existing Items:** Existing items in the child view are exactly the same as in the parent view at the time of configuration. **Item Behavior:*** Branch on Change disabled .	**View Type:** Read-Only Reference **Root Folder:** Selected from parent view **Working Folder:** Usually the same as that of the parent view **Configuration:** Other than floating (Labeled, Promotion State, or As of specified date) **Note:** Changing the item revision attached to a label or changing the label for a promotion state in the parent view may affect items in the child view. These views cannot be rolled back.
Non-derived view Also called Blank branching view **New Items:** New items in child view do *not* appear in the parent view; new items in the parent view do not appear in the child view. **Existing Items:** Existing items in the parent view do *not* appear in the child view. **Item Behavior:** Branch on Change check box disabled.	**View type:** Non-derived **Root Folder:** N/A **Working Folder:** Should be different from that of the parent, to avoid conflicts **Configuration:** N/A

Note You can access Item Behavior by highlighting an item and clicking File (or Folder) > Advanced > Behavior from the menu

If a view is configured to a label, the configuration time for its items is either the time stamp of the label or the time stamp of the item version (if the label was moved). The same is true for a promotion state, which is based on a label. After the creation of a child view, any changes made in the parent view to the label or promotion state on which the child view was based have no effect on the child view.

Creating Additional Views

To create a new view:

1. Display the project view upon which the new view will be based.
2. Select View > New. The *New View Wizard* dialog opens.
3. Select one of the available options from the View Type list box. Depending upon the view currently open, not all view types may be displayed.
   - Branch all (default). All items in the new view start out with their branching behavior set to branch on change. Be sure to use Branch All if you plan to base the new view on a label, promotion state, or a date and time in the parent view. New items and changes to items made in the child view will appear only in the child view. In a floating configuration, new items created in the parent view will appear in the child view. Changes made to an item in the parent view will also be made to the corresponding item in the child view, if that item has not yet branched.
   - Branch none. All items in the new view start out with their branching behavior not set to branch on change, although this option can be changed later. Use Branch None when the new view will have a floating configuration and you want changes to most items in the parent and child views to be reflected in both views, at least at present. If the view has a frozen or fixed configuration, its items will be read-only and none of their data can be changed.

   Note If only a few of your items will branch, select Branch None. If most of your items will end up branched, select Branch All. The branching behavior of individual items can be changed on an item-by-item basis.

   - Reference views. Read/write reference views generally contain a subset of items found in the current configuration of the parent view. They often include only the items of interest to a specific group of workers or an individual. Any change made in a view of this type changes the same item in the parent view. This type of view always floats and its items are not permitted to branch.
   - Read-only Reference views. These views are based upon a specific state of the original view, but are read only. They are typically created for convenience so that revisions of items used in specific product releases can be easily located. For example, a 4.1 Release view might be used to rebuild the 4.1 Release in the future. These views may be float or be frozen. If they float, they may be rolled back; if they do not float, they cannot be rolled back.
   - Non-derived views. As their name implies, these views are not derived from an existing view, even though a parent is selected for them and they display in the view hierarchy. Folders and items are usually added to, moved to, or shared with these views. By default, items placed in these views do not branch, as they do not correspond to items in the parent view.
4. Enter a name and a description of the view in the appropriate text boxes.
5. Click Next. The *New View Wizard: Root Folder* dialog opens.

   From the tree, select the root folder for the new view. Even for a non-derived (blank) view, you must select a root folder.
6. Click Next. The *New View Wizard: Default Working Folder* dialog opens.

   Enter or browse for the name of an appropriate working folder. If you are creating a Reference or Read-Only Reference view, you can use the same working folder as the parent view.

   For a Branch All, Branch None, or Non-Derived view, however, you should always use a working folder that is different from the one used by the parent view. Using the same working folder for both views can cause changes in one view to be overwritten

when files are checked out from the other view. It can also result in incorrect or, at least, misleading file status indicators.

7 At this point, if you are creating a Non-derived view or a Reference view, simply click Finish. It is not necessary to display the *New View Wizard: Configuration* dialog for a Non-derived view because no items from the parent view are included. It is also not necessary to display this dialog for a Reference view because the items have the same configuration as those in the parent view.

8 If you are creating a Branch All, Branch None, or Read-Only Reference view, click Next to display the *New View Wizard: Configuration* dialog.

   a On this dialog, select one of the available configuration options:

   - Floating Configuration

     All the items in the new view will be identical to the corresponding items in the current parent view. Changes to an item in the parent view will be made to the corresponding item in the new view until that item branches, while changes to an item in the new view will be reflected in the parent until the item in the new view branches. (In many cases, the first change to that item will result in branching). New items in the parent view will appear in a branching view. However, new items added to a branching view will appear in the parent view only if the new view is the Branch none type.

   - Labeled Configuration

     Changes to the parent view do not affect the new view. Unless an item is set to Branch on change in the new view, it will be read-only and you cannot change it.

   - Promotion State Configuration

     All the items in the new view will have the promotion state you specify and be rolled back to that state. This option is disabled if the parent view has no promotion states defined for it. Changes to the parent view will not affect the new view. Unless a specific item is set to Branch on change in the new reference view, it will be read-only and you cannot change it.

- Configuration As Of

  The new view will contain only the items that existed at the date and time you specify, and they will be rolled back to just before that date and time. Changes to the parent view will not affect the new view. Unless a specific item is set to Branch on change in the new view, it will be read-only and you cannot change it.

Note: Depending upon the type of view you are creating, not all of these options may be available.

b Click Finish.

Controlling the Branching Behavior of Shared Items

Folders and items can be shared from one view into another if both views belong to the same server configuration. They can also be shared from one folder to another within the same view.

When you share folders and items that can branch, they acquire branching behavior in the new view. Requirements, tasks, and topics do not have branching behavior.

An item's behavior determines whether the item branches on change. A shared item's initial behavior in the new location depends upon the setting of the "Set items shared into view to branch on change" property when the item was shared. This view property appears on the *View Properties* dialog in the root view and in branching views, but it does not appear in reference views. In reference views, the behavior of the shared folders and items that can branch depends on this property's setting in the parent view of the reference view.

After items have been shared into a view, you can change their behavior on an item-by-item basis, but additional changes to the "Set items shared into view to branch on change" property do not change the behavior of the items.

For more information about sharing items, see the *StarTeam User's Guide*.

To control the branching behavior of shared items:

1 Select View > Properties from the menu bar. The *View Properties* dialog appears.

2 Select the Info tab.

3 Select or clear the Set Items Shared into View to Branch on Change check box.

4 Click OK or Apply.

Changing a View's Working Folder

When a view is created, you must specify a location for its working folder. The location of that folder can be changed, but you first need to understand the folder hierarchy and how it is affected by your choice of a working folder.

If a child folder's working folder path is relative to the view's working folder path, changing the view's working folder path will change the child's working folder path as well. The next few paragraphs give you an overview and a couple of examples. For more information, see "Managing Folders" in the *StarTeam User's Guide*.

The working folder for a view's root folder has an absolute path (e.g., "C:\New Product"). The child folder's working folder path depends on how the child folder was created and what changes have been made to it since that time. Usually, the working folder for a child folder is relative to that of its parent. Thus, it is relative to the working folder for the view (that is, the working folder for the view root folder).

For example, suppose that the path to the view's working folder is C:\New Product and that the view's root folder has a child folder named Online Help, with working folder named Online. Then the path to the working folder for Online Help would be C:\New Product\Online.

In this situation, if you change the view's root working folder, you will automatically change the path to the child's working folder. For example, if the view's working folder becomes "C:\Old Product," the working folder for Online Help becomes "C:\Old Product\Online."

To change a view's working folder:

1. Select View > Properties from the menu bar. The *View Properties* dialog appears.
2. Select the Info tab.
3. In the Working Folder group box, select the Default option button.
4. Enter or browse for a new location for the view's working folder.
5. Click OK.

Deleting Views

Be absolutely certain that you want to delete a project view, because you will no longer be able to access any item in the view that is not shared with another project or view. After the view is deleted, it will not be visible in the *Select View* dialog. If other users are connected to the project/view when it is deleted, they will receive a "deleted" message the next time they initiate a view command.

Note Deleting the view does not remove information from the server configuration database. It makes items inaccessible.

To delete the current view:

1. Select View > Delete. A confirmation dialog asks you to confirm the deletion.
2. Click Yes to confirm. A second confirmation dialog requires you to enter the name of the project view.
3. Enter the name (which is case-sensitive) in the View Name text box.
4. Click OK. The open view window for the deleted view closes.

To display another view from this (or another) project, follow the instructions in "Opening an Existing Project" on page 76.

Assigning View Access Rights

This application enables you to assign item-level and view-specific access rights for specific views. In addition to the access rights for a view, you can set item-level and container-level access rights for child folders and specific types of items, because a view is a container for these items. The access rights for these objects apply to all folders and all items of the same type in the view. The access rights dialog for files also has file-specific rights.

You can change these rights on a per-item basis. The access rights dialogs for individual folders and items include only item-level (and, in the case of files, file-specific) rights. For more information regarding project security and access rights, see Chapter 12, "Controlling Access to Objects," on page 189.

Chapter

10

Comparing and Merging Views

Files and other items may need to be merged for a variety of reasons. For example, if both maintenance and new development for a software product need to occur in parallel, a separate view may be created to store each set of source code files. If some defects are fixed in the maintenance view, some items from the maintenance view may need to be merged into the new development view so that the next release of the product will not have the defects.

Here's an outline of this process:

1. From the main baseline view, create a new view for maintenance development.
2. Permit items in the new view to branch and select branch all items.
3. Complete the new view. (Give it a name and select its source folder. Designate a new working folder for the view. Select a view configuration based on a specific date and time, label, or promotion state.).
4. Check all bug fixes into the maintenance view.
5. Do one of the following:
   - Choose View > Compare/Merge to merge changed items back into the main baseline view.
   - Use the viewmerge command at the command line to merge file changes from the maintenance view into the main baseline view.

   You can compare and merge any two project views that are in the same server configuration. Only corresponding items (i.e., items with a common ancestor) can be merged, as they have evolved historically from the same item, even though that item may not be in either of the views you are comparing or merging. Usually you do *not* compare or merge two views that have little in common (that is, with few or no corresponding items).

To merge files, the Windows client provides the View Comparison utility, which has a graphical user interface, and the viewmerge command, which has a command-line interface. Change requests, tasks, or topics can be merged only with the graphical user interface.

Understanding View Comparison and viewmerge Differences

This section explains the differences between the View Comparison utility and the viewmerge command.

The View Comparison utility

- Allows merges from the source view to the target view or from the target to the source.
- Allows the current configuration of a view to be merged with an earlier configuration of itself. The changes resulting from the merge must be applied to the view that has the current configuration.
- Permits two views to be merged using an earlier configuration of one of the views. The changes resulting from the merge must be applied to the view that has the current configuration.
- Allows both file contents and item properties to be compared and merged.
- Never recurses from a folder to its child folders.

The viewmerge command:

- Allows merges only from the source view to the target view.
- Is primarily useful only for merging. Although it compares files, you see the results only when you use Visual Merge to resolve text files with conflicts.
- Merges files based only on their contents.
- Ignores item properties.
- Acts on all files in the view because it recurses through the folders.

When you use recorded merge points, the View Comparison utility always uses the source file (that is, the "merge from" file) from the last recorded merge as the common ancestor. This is the same as using viewmerge with the option usemp. Only the viewmerge command, used with the usemptgt option, allows you to use the target file (that is, the file checked in as a result of the last recorded merge) as the common ancestor.

Using the View Comparison Utility

Regardless of your reasons for merging two views, the steps are similar.

To perform a merge, you usually do the following:

1. Select the views to compare.
2. Select the view configurations to be compared.
3. Decide whether one view, parts of both views, or both views will end up as the merged view.
4. Compare the two views by manually selecting folders or by using Find Differences.
5. Share folders or items that should be in the merged view but are not now there.
6. Delete items from the merged view that are no longer needed.
7. Compare the properties of folders and items in both views and decide what information belongs in the merged view. Merge the folders and items as necessary to combine their properties.
8. Merge text files that are in both views into a single file with no remaining conflicts (lines of text that have changed in two text files or two revisions of the same text file). You may combine their properties, their contents, or both.

9 Save the merged text files in the view that will become the merged view.

10 For binary files, choose the revisions that should end up in the merged view.

Tip You can use the Reconcile command to perform parts or all of steps 5 through 10.

Recording and Tracking Merge Information

When you use the View Comparison utility, the application can record and track merge operations. This feature is useful, especially when you plan to merge the same two views more than once in the same direction.

With merge tracking, the application remembers the last time you merged two corresponding items (one from the source and one from the target view). Regardless of the changes that were made during the merge process, the two items are considered to be identical after the merge. From that time on, it reports that the two items need to be merged only when new changes are made to either or both of the items.

For example, suppose you are using the View Comparison utility to compare a readme.txt file that exists in two different views. The two files have the merge-to-source-or-target icon.

You then merge the corresponding files and check the merged file into the target view. The two revisions that you compared for the merge are considered identical in that direction, even though they are different. The merge-to-source icon replaces the merge-to-source-or-target icon.

When the merge tracking feature is turned off, the application reports that two corresponding items need to be merged when they are not identical. For example, two change requests must have exactly the same property values (except revision numbers) to be identical.

To use merge tracking:

- Launch the View Comparison utility.
- Select Display > Use Merge Tracking For Merge Resolution from the Item or context menu.

To record merge information:

- Launch the View Comparison utility.
- Begin manual or automatic merge operations.
- Select the Record Merge check box.

For more information, see “Merging Item Properties Manually” on page 162, “Merging Files Manually” on page 164, “Merging Item Properties Automatically” on page 166, and “Reconciling Items” on page 166.

Sometimes it is convenient to switch between tracking and not tracking. For example, suppose that you normally use tracking and that you recorded the merge of two corresponding files. Later, you decide that you should have merged them slightly differently. A comparison of the two files confirms that you made a mistake. For example, you may have removed lines that you meant to include.

To rectify this problem:

1 Turn tracking off.

2 Remerge the files.

3 Record the new merge.

4 Turn tracking back on.When you

Displaying a View Comparison Window

When you see the View Comparison utility to compare and merge two views, you display the View Comparison window. hen you see the View Comparison utility to compare and merge two views, you display the View Comparison window.

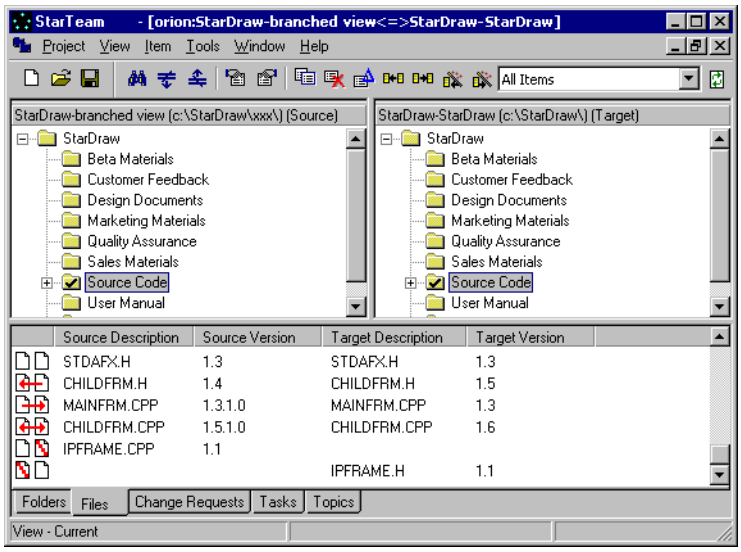

To compare two views:

1 If the first view to be compared or merged is not currently displayed in a view window, display it.

   a Do one of the following:

   - Select Project > Open from the menu bar.
   - Click the Open Project icon on the toolbar.

   The *Open Project Wizard* dialog appears.

   b Click the plus sign in front of the appropriate server configuration to display a list of projects.

Note If you have not yet logged onto the server configuration, you will be asked to do so

   c Select a project, then click Next. The *Open Project Wizard: Select View* dialog appears.

   d Select a view from the hierarchy, then click Finish. This view becomes the source view.

2 (Optional) Configure the view using View > Select Configuration. If you prefer, you can configure the source view after selecting the target view.

3 Next, select View> Compare/Merge from the menu bar.

   a Do one of the following:

   - Select Project > Open from the menu bar.
   - Click the Open Project icon on the toolbar.

   The *Open Project Wizard* dialog appears.

   b Click the plus sign in front of the appropriate server configuration to display a list of projects.

Note If you have not yet logged onto the server configuration, you will be asked to do so.

c Select a project, then click Next. The *Open Project Wizard: Select View* dialog appears.

d From the hierarchy, select a second view, then click Finish. This view becomes the target view.

The two views now appear in the View Comparison window. The view folder hierarchies display side-by-side—the source view in the left pane and the target view in the right pane.

When you select the source pane, the source view configuration displays on the status bar. When you select the target pane, the target view configuration (initially Current) displays on the status bar.

Changing a View or Its Configuration

If you make a mistake or are doing a series of comparisons, you can change to a different source view and/or target view.

You can also change the configuration of either view. Remember that you cannot make changes to any view that has a configuration other than current. This is because you cannot change the past. For example, if you roll the target view back to an earlier time, you can merge only to the source view.

To change to a different source (or target) view:

1 Select the source (or target) pane.

2 Do one of the following:

- From the View or context menu, select Change Project/View to select a view from a different project.
- From the View or context menu, choose Select View to select a different view from the same project.

To roll back a view:

1 Select the source (or target) pane.

2 From the View or context menu, choose Select Configuration to select a label, promotion state, or date and time for the configuration.

To restore a rolled back view:

1 Select the source (or target) pane.

2 From the View or context menu, select Select Configuration to return to the current configuration.

Using the Folder Hierarchies

You can perform some operations on the folder hierarchy in either the source or target pane. For example, you may add a folder, delete a branch, or check folder references or properties.

To add a folder to the view:

1 Select the pane and the folder to be the new folder's parent.

2 Do one of the following:

- Select New Folder from the View or context menu, and choose a folder.
- Drag a folder from Windows Explorer.

To delete a branch from a view folder hierarchy:

1. Select the folder that is the root of the branch to be deleted.
2. Select Delete Folder from the View or context menu.

Note The root folder of a view cannot be deleted.

To check the references for a folder:

1. Select the folder.
2. Select Folder References from the View or context menu.

To review the properties for a folder:

1. Select the folder.
2. Select Folder Properties from the View or context menu.

Comparing Folder Contents Manually

You can manually compare the contents of a folder in the source pane with those of a folder in the target pane. This comparison is only one level deep. That is, you can compare the files in the two folders, but not the files in their child folders. If you want to compare the files in their child folders, you must select the child folders in pairs from the appropriate panes.

To compare folder contents manually:

1. Select a folder from the source view. The application will search for matching folders in the target view. When matching folders are found in the target view, a check mark appears in front of the folders. For example, if the same folder is shared multiple times in the target view, you will see multiple check marks. Remember that no matching folders may exist.
2. Select a folder from the target view. Although the application suggests folders, it will compare any folder that you select.
3. Select a tab from the item (lower) pane to compare the child folders and items in the selected folders.

Now you can review, compare, share, merge, and delete items as explained in later sections of this chapter.

Comparing Child Folders

When you select the Folders tab from the lower pane, you see child folders in the item pane.

Two folders are considered identical when they have the same values for all properties, such as the same description and working folder. Their revision numbers do not have to be identical, and the application does not compare the contents of the folders.

Note An icon that shows two sheets of paper with bent corners identifies identical folders unless the Use Merge Tracking For Merge Resolution option is selected. Then the icon indicates that neither folder's properties have changed since the last recorded merge.

Comparing Files

When you select the Files tab from the lower pane, a list of files displays. Two files are considered identical if:

- Their contents have the same MD5 checksum.
- All their properties have the same values, except revision numbers.

Note An icon that shows two sheets of paper with bent corners identifies identical files unless the Use Merge Tracking For Merge Resolution option has been selected. Then the icon indicates that neither file has changed its properties or contents since the last recorded merge in either direction.

Comparing Items Other than Files

When you select a component tab (Change Requests, Requirements, Tasks, or Topics) from the lower pane, items from that component display.

Any two such items are considered identical if their properties have the same values, except revision numbers.

Requirements, tasks, and topics never branch. Therefore, if one of these items exists in more than one view with a current configuration, changing the item in one view changes it in all of the views. This fact means that source and target views may never have a task or topic that needs merging, unless you reconfigure at least one view.

Note An icon that shows two sheets of paper with bent corners identifies identical items unless the Use Merge Tracking For Merge Resolution option has been selected. Then the icon indicates that neither item has changed since the last recorded merge in either direction.

Controlling the Items That Display

The icons in the item pane represent the relationships between the displayed source and target items, as explained in the following table. The icon on the left in each cell represents the item in the source view and the item on the right represents the corresponding item in the target view.

Icon	Description
No merge needed.	This icon indicates that the item exists in both views and do *not* need to be merged. When the Use Merge Tracking For Merge Resolution option is *not* selected, the corresponding items are identical, which means that they have the same values for all properties, except revision numbers. When the Use Merge Tracking For Merge Resolution option is selected, the items have not changed in either view since the last recorded merge.
Merge to source view (displayed when the Use Merge Tracking For Merge Resolution option is selected).	When an item in the source view can be merged using the corresponding item in the target view, the icon contains a left-pointing arrow. This arrow indicates that only the item in the target view has changed. If the last merge was recorded, the change has occurred since that merge.
Merge to target view (displayed when the Use Merge Tracking For Merge Resolution option is selected).	When an item in the target view can be merged using the corresponding item in the source view, the icon contains a right-pointing arrow. This arrow indicates that only the item in the source view has changed. If the last merge was recorded, the change has occurred since that merge.
Merge to source or target view.	When the corresponding items can be merged to either or both views, the icon contains a double-headed arrow. If the Use Merge Tracking For Merge Resolution option is selected and the last merge was recorded, both items have changed since that merge.

Icon	Description
Item exists only in source.	When an item is missing from the target view, the missing item has a red slash through its icon (the icon on the right). Items with this icon can be shared to the view in which they are missing or deleted from the view in which they exist.
Item exists only in target.	When an item is missing from the source view, the missing item has a red slash through its icon (the icon on the left). Items with this icon can be shared to the view in which they are missing or deleted from the view in which they exist.

Depending on your items and the display options selected for the item pane, not all of these icons may appear.

The application sorts the items in the item pane for you, first by merge status and then by source item descriptor (such as file name or change request number). If all items are displayed, the list will present the items in the following order:

1. Items in both views that *do not* need to be merged.
2. Items in the source view that *do* need to be merged.
3. Items in the target view that *do* need to be merged.
4. Items in both views that *do* need to be merged.
5. Items in the source view only.
6. Items in the target view only.

The icons appear in the same order as the icons in the previous Table. Within each icon group, the items are in ASCII order.

To change the type of item that displays in the item pane:

- Select the appropriate tab at the bottom of the window.

To filter the displayed items:

1. Do one of the following:
   - Select Display from the Item or context menu.
   - Use the Display drop-down list box on the toolbar.
2. Select one of the following filters:
   - All Items

     This filter displays all items that are in the source, target, or both folders.
   - All Items With Differences (default)

     This filter displays all items except those with the two pieces of paper with bent corners icon.
   - All Items Needing Merge

     This filter displays all items that can be merged.
   - Source Items Needing Merge

     This filter displays only the items with changes in the target view that can be merged into the source view.
   - Target Items Needing Merge

     This filter displays only the items with changes in the source view that can be merged into the target view.
   - All Items In Both Views

     This filter displays the items that are in both the source and target folders.

- All Items Not In Both Views

  This filter displays the items that are only one view.

- Items Only in Source View

  This filter displays items in the source folder that are missing from the target folder.

- Items Only in Target View

  This filter displays items in the target folder that are missing from the source folder.

Note The Item > Display menu has the same options as the Item > Select menu. In addition, it has the Use Merge Tracking For Merge Resolution option, which affects both the items that display and the items that can be selected.

Selecting Items from the Item Pane

You can select various subsets of the items currently displayed. Doing this is particularly useful when in a manual comparison of folders.

Depending upon the items displayed, some selections result in the same subset of items. For example, if you display only items with differences, selecting Item > Select > All Items produces the same result as selecting Item > Select > All Items With Differences.

The setting for the Use Merge Tracking For Merge Resolution option affects the meaning of “Needing Merge”. The following procedure assumes that you have selected Item > Display > Use Merge Tracking For Merge Resolution and notes the differences when the option is not selected.

Tip You can use the Display drop-down list box on the toolbar menu to make selections.

To select items by their existence in one or both views:

- For all the items that exist in either or both views, select Select > All Items from the Item or context menu.
- For items that exist in *only* one view, select Select > All Items Not In Both Views from the Item or context menu.
- For the items that exist in both views (not just one view), select Select > All Items In Both Views from the Item or context menu.

To select items by difference:

- For items that differ in any way (i.e., that exist in only one view or differ in both views), select Select > All Items With Differences from the Item or context menu.
- For items that have the same difference:
  - If the items exist *only* in the source folder, select Select > Items Only in Source from the Item or context menu.
  - If the items exist *only* in the target folder, select Select > Items Only in Target from the Item or context menu.
  - If the items exist in both folders and can be merged:
    - For items that have changed in either or both views, select Select > All Items Needing Merge from the Item or context menu.
    - For items that have changed in both views or only in the target view, select Select > Source Items Needing Merge from the Item or context menu.
    - For items that have changed in both views or only in the source view, select Select > Target Items Needing Merge from the Item or context menu.

Note If you have *not* selected Item > Display > Use Merge Tracking For Merge Resolution, the All Items Needing Merge, Source Items Needing Merge, and Target Items Needing Merge options result in the same selection — all items that have changed in either or both views.

Finding Differences

You can move from difference to difference by using the Find Differences command on the View menu. The View Comparison Utility finds the first folder or pair of corresponding folders that meet your criteria for the type of item indicated by your tab selection. Within the item pane, the items within these folders that match your search criteria become selected (provided your display setting enables these items to be displayed).

Critical to the search are:

- The search criteria selected from the *Find Differences* dialog.
- The tab selected from the item pane.

  The search criteria apply to only one type of item at a time and the type is defined your tab selection.
- Whether the source pane or the target pane was active when search options were selected on the *Find Differences* dialog. The active pane is underscored with a color from your color scheme (by default, dark blue).

  You must select the pane for the view to which you want to make changes. Whether it is the source or the target view, it becomes the destination for the merge.

  For example, if the source pane is active (selected), Find Differences locates items in the target that need to be merged to the source.

  If the target pane is active, Find Differences locates items in the source that need to be merged to the target.

  If your goal is to find both types of Find Differences, you must search first in one direction, then in the other. Between searches, you must change the active pane and reset the search options using the *Find Differences* dialog.

  Even if you change the active pane, the Find Next or Find Previous menu commands or toolbar buttons use your original search options and search in the same direction.

Important To locate all items to be merged, regardless of direction (source-to-target or target-to-source), clear merge tracking by selecting Display > Use Merge Tracking For Merge Resolution.

All items to be merged will then show the two-way arrow icon. This information indicates only that the corresponding folders or items differ from one another, rather than which ones have changed since the last merge.

If you choose the wrong display setting, you will not be able to see the items that match your search criteria. For example, if you display items found only in the source folder but search for items found only in the target folder, the items matching your search criteria do not display and therefore cannot be selected.

Caution Remember that although two folders may have the same name, they may not be corresponding folders. Suppose you had a folder named Source Code in the parent view at the time you created a child view that contained all the parent's folders. Then you renamed Source Code in the parent view, calling it Unwanted Code, and created a new Source Code folder in the parent view that was not shared with the child. The View Comparison utility recognizes that the parent's Unwanted Code folder corresponds to the child's Source Code folder and compares the two folders automatically. It will not compare the properties or contents of the two Source Code folders unless you

manually select these two folders from the source and target panes and request a comparison.

To find folders that contain items that match your search criteria:

1. Select an appropriate folder from the source pane.
2. Select an appropriate folder from the target pane.
3. Make either the source pane or the target pane active.
4. Select a tab from the item pane to determine whether the search criteria will apply to folders or items.
5. Carefully select a display setting. Most users select Item > Display > All Items followed by Item > Display > All Items With Differences. This procedure allows them to see all items that match the search criteria, and hide all items that do not match.

   Regardless of the display setting, the application locates the correct folders, but you will almost always want to review the items in the folders matching your criteria.
6. (Optional) To limit the search, select a folder from the inactive pane. Do one of the following:
   - To compare the source view with the target view, select a folder from the target folder hierarchy as the root of the branch to be searched.
   - To compare the target view with the source view, select a folder from the source folder hierarchy as the root of the branch to be searched.

   This step is optional because you are not required to restrict the search within the inactive pane.
7. Indicate which pane will start the comparison by making it the active pane. Select the source or the target pane by:
   - Clicking the pane.
   - Selecting a folder in the pane. The selected folder can also indicate the location from which the search should start.

   If the source pane is active, the application compares the source view with the target view, and vice versa. If neither the source or the target pane is active, it compares the source view with the target view.
8. Do one of the following:
   - Select Find Differences from the View or context menu.
   - Click the Find Differences icon on the toolbar.
   - Press *Ctrl+F*.

   The *Find Differences* dialog appears.

   This dialog's full title is either *Find Differences in Source View* or *Find Differences in Target View*, depending on which pane is active.
9. Select or clear one or more check boxes to indicate the differences in which you are interested.
   - Items That Need To Be Merged

     Select this option to locate corresponding folders that contain items which are different in each folder and need to be merged.
   - Item Only In Source View

     Select this option to locate corresponding folders for which the source folder contains at least one item that is not in the target folder.
   - Item Only In Target View

Select this option to locate the corresponding folders for which the target folder contains at least one item that is not in the source folder.

- No Matching Folder In Source/Target View

  Select this option to locate a folder in the active view that has no corresponding folder in the inactive view. (This option is the only check box unrelated to the tab selected from the item pane.)

10 Select the Root Folder or Selected Folder option button to indicate the location in the active pane from which the search should start.

11 (Optional) Select the Restrict Source/Target View Scope To Selected Folder Branch check box to search only a branch of the folder hierarchy in the inactive pane. The folder selected in the inactive pane is the root of that branch.

12 Click OK. The application finds the first folder or pair of corresponding folders that meet your criteria for the selected type of item. From the item pane, it selects the items within these folders that match your search criteria.

Note If you search only for child folders that have no corresponding folders in the other view, ignore the contents of the item pane. Your search applies only to the folder hierarchy.

If you are using additional search criteria, be careful to distinguish between folders that have *no* corresponding folders (in which case, you ignore the contents of the item pane) and folders that have corresponding folders (in which case, you must be concerned about the contents of the item pane).

13 (Optional) To find the next difference:

- Select Find Next Difference from the View or context menu.
- Click the Find Next Difference icon on the toolbar.
- Press *F3* to find the next difference.

14 (Optional) To find the previous difference:

- Select Find Previous Difference from the View or context menu.
- Click the Find Previous Difference icon on the toolbar.
- Press *Shift+F3* to find the previous difference.

Important Finding the next or the previous difference is based on the criteria previously set for a search, even if you have changed the active pane since that time. Changing the tab at the bottom of the window changes the type of item to which the search is applied, but changing the active window does not change the direction of the search.

Sharing Items

You can share child folders or items if they exist in only one of the selected folders.

When the source view has been set to a fixed configuration (a label, promotion state or date/time), folders and items that are shared to the target view are pinned to that configuration time in the target.

To share the child folders or items:

1 Select one or more items from the item pane. At least one item must have its corresponding item missing from the other view.

If you are selecting requirements, tasks, or topics, you can select only one item. Because these items are in tree form, sharing one item also shares its children.

2 Do one of the following:

- Select Share Items from the Item or context menu.

- Click the Share Items icon on the toolbar. The system displays the following message:

  ```
  Share selected items that doa not exist in both views?
  ```

3 Click OK. A progress bar indicates how close the sharing process is to completion.

Any of the selected items that existed in only one folder are now identical and exist in both folders. Depending on your display setting, they may disappear from the item pane.

Note When folders and items that can branch are shared into a view, they acquire the branching behavior specified for items shared into that view. This branching behavior is a view property setting. See "Controlling the Branching Behavior of Shared Items" on page 146 for details.

Deleting Items

You can delete child folders or items that exist in only one of the selected folders.

To delete child folders or items:

1 Select one or more items from the item pane. At least one item must have its corresponding item missing from the other folder.

  If you are selecting requirements, tasks, or topics, you can select only one item. Because these items are in tree form, deleting one item also deletes its children.

2 Do one of the following:

- Select Delete Items from the Item or context menu.
- Click the Delete Items icon on the toolbar.

  The system displays the following message.

  ```
  Delete selected items that do not exist in both views?
  ```

3 Click OK. A progress bar indicates how close the deletion process is to completion.

When the deletion is complete, the selected items that were in only one folder will not exist in either folder.

Comparing Item Properties

You can compare the properties of any items that exist in both of the selected folders. Remember that items that do not exist in both views cannot be compared.

Comparing files also enables you to compare their contents. See "Comparing Files" on page 162 for more information.

To compare the properties of child folders and items other than files:

1 Select one or more items from the item pane.

2 Do one of the following:

- Select Compare Items from the Item or context menu.
- Click the Compare Items icon on the toolbar.

  If a selected pair of items has the same properties in both the source and target folders, the application displays a message stating that the items are identical.

  If a selected pair of items is different, the *Compare* dialog displays the properties of the items side-by-side. Properties that differ are shown in bold.

3 Use the Previous and Next buttons to move from one pair of items to the next. This process does not wrap from one end to the other.

4 When you are finished, click Close.

Note Custom fields appear in the dialogs that compare properties unless the fields are disabled.

Comparing Files

Comparing files is a special case because the comparison involves both the file revision properties and the file contents.

Files that do *not* exist in both views cannot be compared.

To compare file properties and contents.

1 Select one or more files from the item pane.

2 Do one of the following:

- Select Compare Items from the Item or context menu.
- Click the Compare Items icon on the toolbar.

If a selected pair of files has the same properties and contents in both the source and target folders, the application displays a message stating that the files are identical.

If a selected pair of files is different, the *File Compare* dialog displays the properties of the files side-by-side. Properties that differ are shown in bold.

3 For text files:

a Click Compare Contents to use Visual Diff (or a specified alternate comparison utility) to compare the contents of the tip revisions of the two files.

For more information about Visual Diff, see “Using Visual Merge and Visual Diff” in the *StarTeam User’s Guide*.

b From Visual Diff, click File > Exit to return to the *File Compare* dialog.

4 Use the Previous and Next buttons to move from one pair of items to the next. This process does not wrap from one end to the other.

Merging Item Properties Manually

The properties of child folders or items that exist in both of the selected folders can be merged. For files, you can merge not only the properties but also the contents. See “Merging Files Manually” on page 164.

To merge the properties of child folders or items:

1 Select one or more items from the item pane. At least one item must have an icon indicating that the type of merge you want to perform (to source or to target) is needed.

2 Select Merge Items > To Source or Merge Items > To Target from the Item or context menu. The *Merge* dialog appears. This dialog displays four columns:

- Names of the properties.
- Values of the properties in the source folder or item.
- Values of the properties in the target folder or item.
- Suggested merged properties.

Note Custom fields appear in the dialogs that merge properties unless these fields are disabled.

The merge process involves a comparison of the source and target properties.

- If a property value is the same in both items, that value appears in the Merged column. (This fact indicates that either no changes have taken place or that both the source and target have changed in identical ways.)
- If a property's value is different in both items, the application displays the row in boldface type and places a value in the Merged column based on the following:
  - If the values for a given property were identical at the time of the last recorded merge and only one value has changed since that time, the changed value appears by default in the Merged column.
  - In all other cases, the value that appears in the Merged column by default is the value of the property in the item to which you are merging. For example, if you are merging *to* the source view, the value of that property in the source item appears in the Merged column.

3 Review the properties that appear in boldface type, if any.

4 To change the value in the Merged column to the value in the From or To column:

a Select the property.

b Click Apply From or Apply To.

5 Do one of the following:

- Select the Record Merge check box so that the application remembers the revision of the corresponding item in the other view that involved in the merge.

  If you use this check box during the merge and if you track merges, the merge in the current direction (source or target) becomes unavailable until the corresponding item changes. As a result, two corresponding items may be different but not require a merge.

  When you are not tracking merges, the application indicates that any two corresponding items that are different need to be merged.
- Clear the Record Merge check box if you do not want to record the merge. If you are tracking merges, but do not record this merge, this pair of items continues to be available for a merge in this direction after the current merge is complete.

6 Do one of the following to save the merged folder or item properties as a revision in the source or target:

  Click Apply to remain in the dialog. Using the Previous and Next buttons, move from one pair of items with an appropriate icon to the next. This process does not wrap from one end to the other.If you click Next before Apply, a message box asks for a decision about the current pair of items. Click Yes, No, or Cancel to create a new revision, bypass the current pair, or return to them to finish the merging process.

  If you click Next before Apply, a message box asks for a decision about the current pair of items. Click Yes, No, or Cancel to create a new revision, bypass the current pair, or return to them to finish the merging process.

- Click OK to exit the dialog.

Notes When a new revision is created during a merge, the revision Comment field identifies the item from the other view that was used in the merge. For example, if you merge files, the comment might be:

```
Merge from MAINFRM.CPP, Revision 1.3.1.9.
```

Effects of Merging Change Requests on Workflow

When you merge change requests, you can control which properties end up in which views. However, the workflow for change requests is affected in the following ways:

- If the Last Build Tested and the Addressed In Build fields have build labels as their values (in other words, if these fields are not empty and do not contain the Next Build value), the merged change request retains these values. These values can be changed in the merged view, but only to the names of build labels that exist in that view.
- If the Addressed In Build field contains the Next Build value at the time of the merge, this value is replaced by the next build label created in the original view, not the next build level created in the new view.
- If the Last Build Tested and the Addressed In Build fields in a change request have no values at the time of the merge, their workflow is specific to the view in which they currently reside.

Note The effects on workflow in a merge are the same as those that occur when a change request is moved or initially branches in a branching view.

Merging Files Manually

Both the properties and the contents of files can be merged. When you use the View Comparison utility to create a new file revision, you can change both the file properties and its contents at the same time.

Before performing a merge, ask users to unlock their files (or break the locks yourself), as the application will not create a new revision for a merged file if the file is locked.

To merge two files:

1. Follow steps 1 through 3 in the procedure for Merging Item Properties Manually. (Do not click OK or Apply.)

   The *File Merge* dialog displays the properties for both files and their suggested merged properties. (If no differences exist, this dialog does not appear.)

2. If you plan to merge the contents of the source and target files, do one of the following:
   - Select the Save Merged Item Contents To Server check box to check in the file as part of the merge process.
   - Clear the check box to save the merged contents in the working folder for the “Merge To” location. Doing this enables you to test the merged file before checking it in manually.

   Note Your most recent setting for the Save Merged Item Contents To Server check box becomes a personal option and is stored in the starteam-client-options.xml file.

3. Do one of the following:
   - Select the Record Merge check box to have the application note the revision of the corresponding item in the other view involved in the merge.
     - If you are also tracking merges, the merge in the current direction (source or target) becomes unavailable until the corresponding item changes. Thus, two corresponding items can be different but not need a merge.
     - If you are not tracking merges, the application indicates that any two corresponding items that are different need to be merged.
   - Clear the Record Merge check box if you do not want to record this merge. If you are tracking merges but do not record this merge, this pair of items will continue to be available for a merge in this direction after the merge is complete.

4. Do one of the following:
   - Click Merge Contents to merge the contents of the two files.

- If you are merging text files, follow the directions in “Merging Text Files” on page 165.
- If you are merging binary files, follow the directions in “Merging Binary Files” on page 165.

- Click Apply to save the merged file properties to the source or target view and remain in the dialog.
- Click OK to save the merged file properties to the source or target view and exit the dialog.

Merging Text Files

If you are merging text files, Visual Merge (or a specified alternate merge utility) displays each file. See “Using Visual Merge and Visual Diff” in the *StarTeam User’s Guide* for more information on using Visual Merge.

To merge text files:

1. Follow steps 1 through 3 in “Merging Files Manually” on page 164.
2. Click Merge Contents.
3. Click File > Exit to exit Visual Merge and return to the *File Merge* dialog.
4. Do one of the following:
   - Click Apply to remain in the dialog. Use the Previous and Next buttons to move from one pair of items with an appropriate icon to the next.
   - Click OK to exit the dialog.

Merging Binary Files

With binary files, the merge utility does not display each file.

To merge binary files:

1. Follow steps 1 through 3 in “Merging Files Manually” on page 164.
2. Click Merge Contents. The *Merge File Contents* dialog appears.

   This dialog identifies the two files by name and revision number. The Status field indicates whether the source file, target file, or both have changed (in relation to their common ancestor).
3. Decide which of the files’ contents to use as the merged file by selecting Use Source Contents or Use Target Contents from the Merge Results drop-down list box.

   If both files have changed, you might prefer to compare the two files in the application in which they were created and merge the files together as a single file.
4. Click OK to create the merged file, or click Cancel to avoid creating a merged file at this time. The *File Merge* dialog reopens.
5. To save the merged file’s properties and contents as a revision to the source, target, or both, do one of the following:
   - Click Apply to remain in the dialog. Use the Previous and Next buttons to move from one pair of items with an appropriate icon to the next.
   - Click OK to exit the dialog.

Merging Item Properties Automatically

The auto-merge feature allows you to automatically merge child folders or items. It streamlines the merge process by asking for input only when it finds conflicts (that is, when the same property has changed in both views).

For example, when two files have the same properties or different properties but no conflicting changes to those properties, you will not see the *File Merge* dialog.

Similarly, you see the main window of Visual Merge only if the merged file has conflicts.

To auto-merge child folders or items:

1. Select one or more items from the item pane. At least one item must have an icon indicating that the type of merge you want to perform (to source or to target) is needed.
2. Do one of the following:
   - Select Merge Items > To Source or Merge Items > To Target from the Item or context menu.
   - Click the Merge Items To Source icon or Merge Items To Target icon on the toolbar. The *Merge Options* dialog appears.
3. (Files only) Select or clear the Merge Item Properties and Merge Item Contents check boxes to merge only the file properties, only the file contents, or both.
4. (Files only) If you are merging the file contents, do one of the following:
   - Select the Save Merged Item Contents To Server check box to check the file in as part of the merge process.
   - Clear the check box to save the merged contents in the source or target working folder. Doing this enables you to test the merged file before checking it in manually.

Note Your most recent setting for the Save Merged Item Contents To Server check box becomes a personal option that is stored in the starteam-client-options.xml file.

5. Select the Record Successful Merge check box to store information needed to track merge operations.
6. Select the Use Last Recorded Merge check box to base the merge on the two items being merged and the “merge from” item used in the last recorded merge. Otherwise, the most recent common ancestor will be used to perform a three-way merge. The common ancestor is a historical revision that has the same revision number in both views.
7. Click OK.The *Merge* dialog, merge utility, and *Merge File Contents* dialog appear so that conflicts can be resolved. See “Merging Item Properties Manually” on page 162 and “Merging Files Manually” on page 164 for details.

If no conflicts exist, you will see only the *Merge In Progress* dialog, which displays a progress bar.

Tip The *Merge* dialog contains a Skip button during the auto-merge process so that you can bypass a pair of items instead of merging them.

Reconciling Items

The reconcile process is an automated process for sharing, deleting, and merging the folders or items displayed in the item pane. All of these processes can be performed manually, using the reconcile process save you several steps if all selected items with the same icon are to be treated identically.

To reconcile child folders or items:

1. Select one or more items from the item pane. At least one of these items must have a difference, as indicated by its icon.
2. Do one of the following:
   - Select Reconcile > Source View or Reconcile > Target View from the Item or context menu.
   - Click the Reconcile Source icon or Reconcile Target icon on the toolbar.

   The *Reconcile* dialog appears.
3. Indicate the action to take with selected items that exist only in the source view. From the For Items Only In Source View group box, select the Do Nothing, Share Items To Target View, or Delete Items From Source View option button.

   One of the first two group boxes (For Items Only In Source View or For Items Only In Target View) will be disabled, depending on the type of merge you are performing.
4. Indicate what should be done with selected items that exist only in the target view. From the For Items Only In Target View group box, select the Do Nothing, Share Items To Source View, or Delete Items From Target View option button.

Note One of the first two group boxes (For Items Only In Source View or For Items Only In Target View) will be disabled, depending on the type of merge you are performing.

5. Indicate what should be done with selected items that exist in both views but are different. From the For Items Needing Merge group box, select the Do Nothing, Merge Items, or Auto-Merge Items option button.
6. Click OK. Depending on your selections, various *In Progress* dialogs open and display progress bars.
7. If you selected the Auto-Merge Items option button, the *Merge Options* dialog appears.
   a. (Files only) Select or clear the Merge Item Properties and Merge Item Contents check boxes to merge only file properties, only file contents, or both.
   b. (Files only) Do one of the following:
      - Select the Save Merged Item Contents To Server check box to check the file in as part of the merge process.
      - Clear the check box to save the merged contents in the working folder for the “Merge To” location. Doing this enables you to test the merged file before checking it in manually.

Note Your most recent setting for the Save Merged Item Contents To Server check box becomes a personal option that is stored in the starteam-client-options.xml file.

   c. Select the Record Successful Merge check box to store information needed to track merge operations.
   d. Select the Use Last Recorded Merge check box to base the merge on the two items being merged and the “merge from” item used in the last recorded merge. Otherwise, the most recent common ancestor will be used to perform a three-way merge. The common ancestor is a historical revision that has the same revision number in both views.
   e. Click OK.

If you selected the Merge Items or Auto-Merge Items option button, the *Merge in Progress* dialog appears and is interrupted by the other dialogs that the application uses during a merge or auto-merge process. For more details, see “Merging Item

Properties Manually” on page 162, “Merging Files Manually” on page 164, or “Merging Item Properties Automatically” on page 166.

Reviewing Item References

When you review item references, you see all the references for the selected item or folder, even if the reference is in a view that is not in the View Comparison window.

To check references for a child folder or item:

1. Select a folder or item from the lower pane.
2. Select Item > Item References. The <item> References dialog appears. Arrows point to the references. Depending on whether you selected an item that appears in one or both views, you will see one or two arrows.

   The arrows for the source and target files are slightly different:

   - The arrow pointing outward is for the source item reference.
   - The arrow pointing inward is for the target item reference.

Reviewing Item Properties

You can review the properties of any item displayed in the item pane.

To review properties for a child folder or item:

1. Select one or more items from the item pane.
2. Do one of the following:

   - Select Source Item Properties from the Item or context menu.
   - Click the Source Item Properties icon on the toolbar.
   - Select Target Item Properties from the Item or context menu.
   - Click the Target Item Properties icon on the toolbar.

   The *Properties* dialog displays the properties of the first selected item in the source or target folder, depending on your selection.
3. Click Previous or Next to move to the *Properties* dialog for the preceding or next selected item.

Tip To see the properties for a single item, double-click the descriptor (for example, a file name or change request number) in the Source Description or Target Description column of the item pane.

Refreshing the Panes

You can refresh individual panes or the entire View Comparison window. Activating a pane (for example, by selecting an item from the pane or by right-clicking the pane to use the context menu), always refreshes it. The menu commands described below also refresh the panes.

To refresh the item pane:

- Do one of the following:
  - Select Refresh from the Item or context menu.
  - Click the Refresh Item List icon on the toolbar.

To refresh the source or target pane:

- Select Refresh from the View or context menu. Updating the source or target pane also often updates the item pane.

To refresh all the panes:

- Select Window > Refresh.

Merging Views from the Command Line

As noted earlier, the viewmerge command is an alternative method of merging merge file contents and sharing or deleting items. This command ignores file properties (such as the file description) and merges files based only on their contents.

By default, viewmerge compares all the files in the source view with all the corresponding files in the target view. Using the -sourcepath and -targetpath options, you can compare any branch in the source view folder hierarchy with any branch in the target view folder hierarchy. Using the -norecursion option, you can even limit the comparison to a single folder in each view.

By default, this command merges only text files. During a merge, the specified merge utility displays and compares text files that contain conflicts. Using the the utility, you can resolve the conflicts before viewmerge checks in the merged file and goes to the next pair of files.

If the merge output generated by the utility matches the target file's current contents, the target file is not checked in, because there is no reason to do so.. If you use the -usemp or -usemptgt option and the source file has not changed since the last merge, no merge occurs, because nothing needs to be added to the target file. If you use the -recordmp option, a merge point is created in either of these cases.

If at least one file in the pair to be merged is a binary file, viewmerge sends an error message to standard error.

The -noconflicts, -usesource, and -usetarget options enable you to deviate from a default merge. All of these options stop viewmerge from using the merge utility and provide other ways to handle files that have conflicts. The -usesource or -usetarget options also provide methods of managing binary files.

With the -noconflicts option, viewmerge merges only text files that have no conflicts, bypassing text files that have conflicts and binary files. An error message is sent to standard error for these files.

With the -usesource text and -usetarget text options, viewmerge allows you to manage text files that have conflicts by using the existing tip revision from the source or the target view for the merged file. With the -usesource binary and -usetarget binary options, viewmerge allows you to manage binary files with differences by using the existing tip revision from the source or the target view for the merged file.

Notes

- When a new revision is created using viewmerge, the revision's Comment field provides merge information. The Comment identifies the item from the other view that was used in the merge. For example, if you merge files, the comment might be:

  ```
  Merge from MAINFRM.CPP, Revision 1.3.1.9
  ```

- If a file is locked by another user in the view to which you are merging , the application does not create a new revision with the merged file. Therefore, you should ask users to unlock their files or break their file locks yourself before performing the merge.

The viewmerge syntax uses the following conventions:

Convention	Description
Bold	Information that you must use exactly as shown (if you use it).
Italics	Information that you replace with the names of your files, child folders, etc.

Convention	Description
[ ]	Square brackets surround optional syntax.
|	A vertical bar separates mutually exclusive choices. You select only one of the choices.

This table illustrates the syntax you would use for the command:

Command	Syntax
viewmerge	**-s** "userName:password@hostName:endpoint" **[-pwdfile** "filePath"**] [-cmp]** **[-encrypt** "encryptionType"**]** **[-nologo] [-?]** **-source** "projectName/viewName" **[-sourcepath** "folderPath"**]** **-target** "projectName/viewName" **[-targetpath** "folderPath"**]** **[-srclabel** "viewLabel" **|** **-srcstate** "promotionState" **|** **-srctime** "dateAndTime"**]** **[-norecursion] [-usecontents]** **[-nomerge |** **-noconflicts [locktype [source | target] |** **-outofdate -recordmp -usemp|** **-usesource [text | binary]|** **-usetarget [text | binary]]** **[-noci] [-dry] [-b]** **[-recordmp] [-usemp| -usemptgt]** **[-si source | target] [-sf source | target]** **[-di source | target] [-df source | target]** **[-itemtype FILE] [-itemtype CR]** **[-itemtype REQ] [-itemtype TASK]** **[-itemtype TOPIC]**

viewmerge Options	Description
-?	At the command line, lists the command's syntax and a description of each option.
-b	When two lines of text files are compared, ignores trailing white space and treats all other strings of white space as equal in length. For example, the following lines are equivalent: " hi mom " " hi mom" This option is meaningless if you use -nomerge.
-cmp	Compresses all the data sent between the workstation and the server and decompresses the data when it arrives. Without this option, the default is no compression. Compression is most useful and appropriate when the client and server communicate over a slow connection. You must consider whether the time spent compressing and decompressing data is better than the longer transfer time of uncompressed data sent over the slow connection.

viewmerge Options	Description
`-df`	Deletes folders from the specified view if they do not exist in the unspecified view.[1] The full syntax is: `-df source | target` You indicate which view will change as a result of using this option. For example, `-df target` deletes folders from the target view that do not exist in the source view. Deleting a folder deletes its entire contents. The command deletes the child folders and items in the deleted folder, recursively. This action includes files, change requests, requirements, topics, and tasks. You may use this option more than once in a command if you want to change both the source and the target. See also -sf.
`-di`	.Deletes items of the specified types from the specified view if they do not exist in the unspecified view. This option applies to items in folders that exist in both views. It includes files, change requests, requirements, topics, and tasks. The full syntax is: **`-di source | target`** You must indicate which view will change as a result of using this option. For example, `-di target` deletes items from the target view that do not exist in the source view. You may use this option more than once in a command if you want to change both the source and the target. See also -si.
`-dry`	Performs a “dry run” in which items are compared but not merged. With this option, the command lists the following (in addition to the usual viewmerge output): ■ The number of files that can be merged because they have content differences. These differences exist since the last recorded merge if -usemp or -usemptgt is used. ■ The number of files that can be merged that also contain conflicts. ■ The number of folders and files found in one view and not in the other. If you use -dry with -nomerge, the command checks for and lists only the number of folders and files found in one view and not in the other. It supplies no merge information. If you use -dry with -noci, -noconflicts, -usesource, or -usetarget, -dry takes precedence. If you use -recordmp with -dry, viewmerge ignores the -recordmp option.
`-encrypt`	Encrypts all the data sent between the workstation and the server and unencrypts it when the data arrives. Without this option, no encryption takes place. Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines. The full syntax is: **`-encrypt`** “*encryptionType*” The types of encryption are: **RC4** RSA R4 stream cipher (fast) **RC2_ECB** RSA R2 block cipher (Electronic Codebook) **RC2_CBC** RSA R2 block cipher (Cipher Block Chaining) **RC2_CFB** RSA R2 block cipher (Cipher Feedback) These encryption types are ordered from fastest to slowest. Each of the encryption types is slower and safer than the one preceding it.

viewmerge Options	Description
`-itemtype`	Indicates which item types are processed for item-sharing and item deleting operations (specified with the -si, and -di options). When -itemtype is not used, these operations apply to all item types. When -itemtype is used, it must be used once for each type that will be shared or deleted. The full syntax is: **[-itemtype FILE] [-itemtype CR] [-itemtype REQ] [-itemtype TASK] [-itemtype TOPIC]** FILE, CR, REQ, TASK, and TOPIC stand for files, change requests, requirements, tasks, and topics, respectively. If no item-sharing or item-deleting options are used in the command, the itemtype options are ignored.
`-noci`	Does not check in merged files to target view. Instead, the merged files are saved as working files in the target view's working folders. You can use a merge utility to resolve files with conflicts unless -noconflicts is also used. This option is meaningless if you use -nomerge. If you use -recordmp with -noci, viewmerge ignores the -recordmp option.
`-noconflicts`	Merges only text files that have no conflicts. If you use this option, viewmerge does not display the merge utility when the text files to be merged have conflicts. Instead, viewmerge sends error messages to standard error for these files (in addition to the error messages it sends for binary files). Files with conflicts are left as they were prior to running the viewmerge command unless you specify a locktype. Then the files with conflicts are locked, using the user ID of the person running the viewmerge command. The files can be merged later, when someone is available to review them in the merge utility. Using -noconflicts enables you to run viewmerge without interruption and resolve file conflicts later. The full syntax is: **-noconflicts [*locktype* [source | target]** where *locktype* can be either an exclusive lock (**elock**) or a non-exclusive lock (**nlock**). If *locktype* is not used, no files are locked. When you specify a locktype, you can lock: ■ Only the source file by using *locktype* **source**. ■ Only the target file by using *locktype* **target**. ■ Both the source and target files by using *locktype* alone. Even when -recordmp is used, the viewmerge command does not record a merge point for files that have conflicts. It just moves to the next pair of files. This option is meaningless if you use -nomerge. If you use -noci with -noconflicts, nothing is checked in. Otherwise, viewmerge checks in the merged files for files that have no conflicts. In addition, when you use these options together, the -recordmp option is ignored.
`-nologo`	Suppresses the copyright notice.
`-nomerge`	Does not merge files. Use this option when you want to share or delete items without merging files. Without this option, files in the source view are merged with their corresponding files in the target view, if a merge is needed. If you use -nomerge, all other merge options, such as -noconflicts, -outofdate, and -recordmp, are ignored. If you use -nomerge with -dry, the command lists the number of missing folders and items.
`-norecursion`	Prevents recursion to child folders so that only the starting folders in the source and target views are compared. These folders may be specified with -sourcepath and -targetpath. If they are not specified, the default is the root folder for each view.

viewmerge Options	Description
`-outofdate`	If the source file has changed but the target file has not, this option treats the target file as an out-of-date file and checks in the source file as its replacement. This option works with both binary and text files. It ignores files that have changed in both views or only in the target. When this option is used with -recordmp and -usemp, it recognizes files as out-of-date based on the previous merge in the same direction. Otherwise, it bases the decision on the common ancestor. Borland recommends using -recordmp and -usemp.
`-pwdfile`	Specifies the complete path to a file that stores the user's password. This option overrides the password used as part of the -s option. It prevents others from seeing the user's password on the command line. The full syntax is: **`-pwdfile "`***filePath***`"`** If you are inputting the password from a password file, the -s option syntax changes to remove the password but not the colon (:) between the user's name and the password. **-s** "*userName*:*@hostName***:***endpoint*" The password file should contain only the password. Leading and trailing whitespace is ignored.
`-recordmp`	If you use the -recordmp option, a merge point is recorded for each pair of files if: ■ A merged file is created from the pair. ■ The file generated by the merge utility matches the current contents of the target file, so no target file is checked in. ■ No merge takes place because the -usemp or -usemptgt option has been used and the source file has not changed since the last merge in this direction. ■ You use the -usesource or -usetarget option. When a pair of files is merged using the -usemp or -usemptgt option, the source file from the last recorded merge is used as the common ancestor of the files. If no previously recorded merge occurred, the common ancestor is used. This option is meaningless if you use -nomerge, -noci, or -dry.
`-s`	Specifies the user's name and password and the server's address and endpoint as follows: *userName:password@hostName***:***endpoint* For example: `-s "JMarsh:password@orion:49201"` If the user name is omitted, the current user name is used. The user name in the example is "JMarsh" If the password is omitted, the user is prompted to enter the password. The password in the example is "password" If the host name is omitted, the default is localhost. The host name in the example is "orion". The default endpoint (port) for TCP/IP is 49201. The endpoint in the example is 49201.

viewmerge Options	Description
`-sf`	Shares folders into the specified view from the unspecified view if they do not already exist in the specified view. The full syntax is: **`-sf source	target`** You indicate which view will change as a result of using this option. For example, `-sf target` shares folders into the target view that exist only in the source view. Sharing a folder shares its entire contents. The command shares the child folders and items in the shared folder, recursively. This includes change requests, requirements, topics, and tasks, as well as files. You may use this option more than once in a command if you want to change both the source and the target. See also -df.
`-si`	Shares items of the specified types into the specified view from the unspecified view if they do not already exist in the specified view. This option applies to items in folders that exist in both views. It includes change requests, requirements, topics, and tasks, as well as files. The full syntax is: **`-si source	target`** You must indicate which view will change as a result of using this option. For example, `-si target` shares items that exist only in the source view into the target view. You may use this option more than once in a command if you want to change both the source and the target. See also -di.

`-source`	Required. Specifies the source view.
	The full syntax is:
	`-source` "*projectName/viewName*"
	For example:
	`-source "Big Project/Child View"`
	The project name in the example is Big Project. The view name in the example is Child View, a child of the root view.
	The project name is always required.
	If the view name is omitted, the root (default) view is used. The root view starts out with the same name as the project, but its name can be changed. If two child views have the same name, you can eliminate the ambiguity by replacing *viewName* in the view to be used with the hierarchy found in the root view. Otherwise, the command will use the first view it locates with the name.
	The following example provides both a project name (Big Project) and the path from the root view (also named Big Project) to a view that has the same name as another view:
	`-`**`source`** `"Big Project/Big Project/Branch 8/Child View`

`-srctime`	You can configure (or roll back) the source view to a specific point in time. The full syntax is: **`-srctime`** "*dateAndTime*" Examples include: "12/29/01 10:52 AM" "December 29, 2001 10:52:00 AM" "Monday, December 29, 2001 10:52:00 oclock AM" This option is mutually exclusive with -srclabel and -srcstate. When no option is used, the view configuration is the current configuration.

viewmerge Options	Description
`-sourcepath`	Specifies the path to the starting folder in the source view. The default is the root folder of the source view. The full syntax is: **`-sourcepath`** "*folderPath*" For example: -sourcepath `"BigProject/SourceCode"`
`-srclabel`	You can configure (or roll back) the source view to a specified view label. Then you can compare and merge the view as it was at that point in time. The full syntax is: `-srclabel` "*viewLabel*" This option is mutually exclusive with -srcstate and -srctime. When no option is used, the view configuration is the current configuration.
`-srcstate`	You can configure (or roll back) the source view to the view label assigned to the specified promotion state. Then you can compare and merge the view as it was at that point in time. The full syntax is: `-srcstate` "*promotionState*" This option is mutually exclusive with -srclabel and -srctime. When no option is used, the view's configuration is the current configuration.
`-srctime`	You can configure (or roll back) the source view to specific point in time. The full syntax is: `-srctime` "*dateAndTime*" Examples include: "12/29/01 10:52 AM" "December 29, 2001 10:52:00 AM" "Monday, December 29, 2001 10:52:00 oclock AM" This option is mutually exclusive with -srclabel and -srcstate. When no option is used, the view's configuration is the current configuration.
`-target`	Required. Specifies the target view. The full syntax is: **`-target`** "*projectName/viewName*" See -source above for details.
`-targetpath`	Specifies the path to the starting folder in the target view. The default is the target view's root folder. The full syntax is: **`-targetpath`** "*folderPath*" For example: -targetpath `"BigProject/SourceCode"`
`-usecontents`	Bases the merging of text files on the contents of the source and target files, rather than their MD5 checksum (which is the default basis for text file comparison).
`-usemp`	Bases the merge on the source file from the last recorded merge as well as the pair of files being merged. (The mp stands for merge point.) Without -usemp or -usemptgt, the most recent common ancestor is used to perform the three-way merge. The common ancestor is a historical revision with the same revision number in both views. This option is meaningless if you use -nomerge.
`-usemptgt`	Bases the merge on the results from the last recorded merge as well as the pair of files being merged. The results from the last recorded merge were checked in to the target view. (The mp stands for merge point and the tgt stands for target.) Without -usemp or -usemptgt, the most recent common ancestor is used to perform the three-way merge. The common ancestor is a historical revision with the same revision number in both views. This option is meaningless if you use -nomerge.

viewmerge Options	Description
`-usesource`	"Merges" all files of the specified type, text or binary. When no type is specified, all files are "merged". If no type is specified, all files are "merged." ("Merges" is in quotation marks because the source file may be used as the merged file.) With this option, viewmerge does not display the merge utility when text files have conflicts. The full syntax is: **`-usesource`** [**`text`** | **`binary`**] The source file is substituted for the merged file and checked in to the target view if: ■ At least one of the files in a pair of files to be merged is binary. ■ A pair of text files to be merged contains conflicts. The viewmerge command records the merge point if -recordmp is used and moves to the next pair of files. This option is meaningless if you use -nomerge.
`-usetarget`	"Merges" all files of the specified type, text or binary.If no type is specified, all files are "merged". ("Merges" is in quotation marks because the target file may be used as the merged file.) With this option, viewmerge does not display the merge utility when text files have conflicts. The full syntax is: **`-usetarget`** [**`text`** | **`binary`**] The target file is substituted for the merged file and checked in to the target view if: ■ At least one of the files in a pair of files to be merged is binary. ■ A pair of text files to be merged contains conflicts. However, it is never necessary to check in the target file, as it already is the tip revision. The viewmerge command records the merge point if -recordmp is used and moves on the next pair of files. The differences between this option and the -noconflicts option are: ■ If you use -recordmp, that option applies to all pairs of files, not just text files without conflicts. ■ No error messages are sent to standard error for binary files or for text files that have conflicts This option is meaningless if you use -nomerge.

[1] Be aware of the effects of using -sf and -df or -si or -di together. For example, if you use -sf source with -df source, -sf source will add folders to the source that were only in the target, while -df source will remove folders from the source that were not in the target. These two options work with different sets of folders. Using -sf source and -df target together moves folders that were not in the source, but were in the target, to the source. Then the folders are removed from the target.

Logging Information

During the viewmerge process, data appears on the screen explaining what is happening. This data has been sent to standard output or to standard error. If you wish, you can have this information logged instead of displayed.

For example, on the Windows client , if you end the viewmerge command with "1>2& viewmerge.log," the data that is sent to standard output and standard error appears in viewmerge.log. Ending the command with "> viewmerge.log" logs only the data from standard output.

viewmerge Examples

The following viewmerge command merges the contents of files that need to be merged from the source view to the target view. The source view is the root view of the project named Big Project, while the target view is the view named Child View in the project named Big Project.

Because the command includes the -usemp option, the command determines whether two corresponding files need to be merged by whether the item in the source view (in this case, Child View) has changed since the last recorded merge.

Because the command includes the -recordmp option, the merge will be recorded as a merge point, and the merged file will be used the next time the same pair of files is merged with this command, the same source and target views, and the -usemp option.

In addition to performing the merge operations, this command shares folders and items from the source view into the target view. The shared items and folders are those that exist in the source view but are missing from the target view.

Because the command ends in "> viewmerge.log", the data sent to standard output is copied to the viewmerge.log file.

```
viewmerge -s "JMarsh:password@orion:49201" -source "BigProject" -target "Big
Project/Child View" -recordmp -b -usemp -si target -sf target > viewmerge.log
```

The next viewmerge command merges the contents of files that need to be merged from the source view to the target view. However, because neither the -usemp nor -usemptgt option is used, recorded merge points are ignored. Because the -recordmp option is not used, this merge does not record a merge point for future merges.

If the source file and the target file are different, the two files are merged. This command can be run with no supervision because the -noconflicts option bypasses files that have conflicts.

```
viewmerge -s "JMarsh:password@orion:49201" -source "BigProject" -target "Big
Project/Child View" -noconflicts -b > viewmerge.log
```

Without -noconflicts, viewmerge displays the merge utility for files that have conflicts so that you can resolve the conflicts.

If the person running the viewmerge command cannot resolve the conflicts, using -noconflicts enables others to merge the files with conflicts at a later time by using the merge utility from the command line or from the View Comparison window.

The next viewmerge command is similar to the previous one. However, instead of ignoring files with conflicts, this command checks the source file into the target view as though it were the merged file.

```
viewmerge -s "JMarsh:password@orion:49201" -source "BigProject" -target "Big
Project/Child View" -usesource > viewmerge.log
```

In the next example of the viewmerge command, files in branches of the views' folder hierarchies are compared by using -sourcepath and -targetpath.

```
viewmerge -s "JMarsh:password@orion:49201" -source "BigProject" -sourcepath
"Big Project\Source Code" -target "Big Project/Child View" -targetpath "Big
Project\Source Code" > viewmerge.log
```

The next viewmerge command is similar to the previous one. However, the -norecursion option limits the recursion to exactly one folder (Source Code) and ignores any child folders.

```
viewmerge -s "JMarsh:password@orion:49201" -source "BigProject" -sourcepath
"Big Project\Source Code" -target "Big Project/Child View" -targetpath "Big
Project\Source Code" -norecursion > viewmerge.log
```

Troubleshooting

Version and Build Numbers

If you encounter problems with the viewmerge command, you may need to know the version and build number for your copy of viewmerge. The first few lines sent to standard output contain that information. The following three lines are a sample:

```
StarTeam 6.x Command Line View Merge Utility
Copyright (c) Borland Software Corporation
Build Number 5.x.xxx
```

Empty Files

When a merge results in an empty file, despite the fact that the files to be merged were not empty, an error message is sent to standard error. For example, if the common ancestor file has no end-of-line characters, the merge utility may think the merged file should be empty.

Changed Names of Folders and Files

Folder and file name changes are not propagated. For example, suppose the source folder f1 is renamed as f2 and then compared to the target folder named f1. The two folders will match because they have the same root object ID, but the source folder's name change will not propagate to the target folder.

Input/Output Errors

No checks are made for I/O errors. For example, if a file cannot be checked-out, copied, or opened for comparison, no error is reported.

Chapter

11

Customizing Item Properties

Enterprise Advantage and Enterprise customers can customize the repository by adding user-defined property fields to files, change requests, requirements, topics, and tasks.

The properties of some existing items can also be changed. For example, Priority is a change request field that can be customized. Instead of using the default values (Yes and No), your company can change the item properties so that change requests are prioritized on a scale from 1 (high) to 10 (low).

Important When you customize the database, you should use the stcmd command line to exclusive lock the server. For information, see the StarTeam User's Guide.

The application adds new property fields and modified property fields to the database used by your current server configuration. These fields are available in the same places as other item property fields.

To display user-defined fields on the upper pane:

a. Select Filters > Show Fields from the appropriate menu, or right-click to display the context menu and select Filters > Show Fields. The *Show Fields* dialog appears.

b. Select the fields you want to display from the Available Fields list and click Add.

c. Arrange the fields in the Show These Fields In This Order list by dragging and dropping the field names.

d. Click OK when you are finished.

e. The fields you select appear in the upper pane.

Note: Only 60 columns can be displayed in the upper pane at a time. Therefore, if adding custom fields makes over 60 fields available, you will not be able to display them all in the upper pane. Also, remember that databases have limitations which limit the number of custom fields you can create.

You can also see a list of the custom fields for a particular type of item on the Custom tab of the item *Properties* dialog.

Creating a User-Defined Property Field

The following table lists the data types you can use when creating a user-defined field.

Table 11.1 User-Defined Property Field Types

Field Type	Description
Real	Indicates that the field is numeric and can contain a decimal point. The range is 1.7E +/- 308 (15 digits); however, the range may be limited by your database. For example, Microsoft SQL Server ignores all but the first 32 characters.
Enumerated	Indicates that the field will contain a value from a set of enumerated values. For example, the enumerated values for a field named “Traffic light” would be red, yellow, and green.
Integer	Indicates that the field contains a whole number. The range is limited to 9 digits to ensure that you can migrate from database to database successfully.
Text	Indicates that the field is a string of the length that you specify.
Date/Time	Indicates that the field will contain a date and time.
User ID	Indicates that the field will contain a user’s name from the set of values in the list of users.

To create a user-defined property:

1 Select Advanced > Customize from one of the item menus, such as the File or Change Request menu.

The *Customize* dialog lists the fields that can be customized:

- A field with a pencil in the center of the icon is an application field. It is always an enumerated type and is fully customizable. You can add, disable, rename, or reorder the enumerated type’s values.
- A user-defined field has an icon displaying the head and shoulders of a person. It can be one of several types and is fully customizable. If it is disabled, the icon becomes greyed-out.
- A field with an icon containing a small yellow lock in the lower left corner is a restricted enumerated type. Only the names the application displays for the enumerated type’s values can be changed. Usually these fields have workflow characteristics that cannot be altered, such as the change request field named Status. The value of this field affects how the change request is processed and what its next value can be.

2 Click Add. The *Add Field* dialog appears.

3 In the Field Name text box, enter the name to be used by the database for the user-defined field. Use only alphanumeric characters and no spaces in this name.

The name should be less than 31 ASCII characters (including the Usr_) and should not contain any of the following characters, which are not accepted by one or more of the databases that the application supports:

```
= \\.^$@,;!:#*&<>?-//%|[](())+"
```

Be careful about selecting this name, as it cannot be changed after you click OK.

4 Enter the name that the application shows to users in the Display Name text box.

The Server allows you to use the same display name for more than one custom field. Depending on your use of custom fields, doing this may not be good practice.

5 From the Type list box, select a type and follow the steps in the procedure for that type.

- For enumerated types, see “Creating an Enumerated Type” on page 181.

- For integer or real types, see “Creating a Numeric Type” on page 181.
- For text types, see “Creating a Text Type” on page 182.
- For date/time types, see “Creating a Date/Time Type” on page 182.
- For user ID types, see “Creating a User ID Type” on page 183.

Enumerated, integer, real, and date/time fields have default values so they always have a value, although the end user can change it. Text and user ID fields do not have default values, but you can customize the fields to require user input.

For details about translating the display names and enumerated values for custom fields, see “Translating Display Names and Enumerated Values” on page 186.

Creating an Enumerated Type

This procedure is a continuation of “Creating a User-Defined Property Field” on page 180.

In the *Query* and other dialogs, the order in which the enumerated values appear is the order in which they appear in the *Add Field* or *Modify Field* dialog. This order can be code order or even alphabetic order, but only if you arrange them that way. To rearrange the values, you can use drag-and-drop.

To create an enumerated type:

1. Follow the steps listed under the heading “Creating a User-Defined Property Field” on page 180 and then select Enumerated from the Type text box.
2. Click Add to enter the first value for this enumerated type. The *Add Value* dialog appears. The application reserves the numeric codes from 0 to 100, so this dialog shows numbers starting with 101.
3. Use the displayed code or type another in the Code text box.
4. Enter the name for the enumerated value in the Name text box.
5. Click OK.
6. Repeat Steps 2 through 5 to add additional values.
7. When you have finished adding values, use drag-and-drop to arrange the values in the Possible Values list box.
8. Select one of the enumerated values as a default from the Default Value list box. This default automatically becomes field value for the field in all existing items for which the field is a property. It also becomes the default value for newly created items. You can change the value manually on the Custom tab of the item *Property* dialog.

   Many users create a value of <none> to use as the default and to indicate that no value is really being set. They often place this value first in the enumerated list so that it will sort to the top or bottom of the upper pane when the column for the custom field is sorted.

   If you are entering a date or time, use the format for your locale.
9. Click OK.

Caution Before you click OK, verify that the codes you have selected are correct. They cannot be changed after you exit the *Add Field* dialog.

Creating a Numeric Type

This procedure is a continuation of “Creating a User-Defined Property Field” on page 180. Numeric fields can have either Integer or Real as their types.

To create an integer or real type:

1. Follow the steps listed under the heading "Creating a User-Defined Property Field" on page 180 and then select Integer or Real from the Type text box.
2. Enter a default integer or real value in the Default Value text box. This default automatically becomes the value for this field in all existing items for which the field is a property. It also becomes the default value for newly created items. You can change the value manually on the Custom tab of the item *Property* dialog.
3. Click OK.

Creating a Text Type

This procedure is a continuation of "Creating a User-Defined Property Field" on page 180. Remember that you cannot set a default value for a text field, because most text fields contain unique data.

To create a text type:

1. Follow the steps listed under the heading "Creating a User-Defined Property Field" on page 180 and then select Text from the Type text box.
2. Use the default maximum length (255 characters) or enter a number of characters from 2 to 20,000 in the Length text box. You cannot change this length after you click OK, so be sure it is adequate for the field.
3. Select the Input Required check box if you want to make this a required field.

Note Required fields are meaningless when you add, check in, and check out files. Because they affect only the *File Properties* dialog, they are rarely used for files.

4. Click OK.

Creating a Date/Time Type

This procedure is a continuation of "Creating a User-Defined Property Field" on page 180. Date and time formats vary from region to region and company to company. For example, some companies use the date format mm/dd/yy, while others use the date format dd/mm/yy. The application displays dates and times using the formats defined in the *Windows Regional Settings* dialog.

To create a date or time type:

1. Follow the steps listed under the heading "Creating a User-Defined Property Field" on page 180, and then select Date/Time from the Type text box.
2. (Optional) Enter a default date or date and time in the Default Value text box. You cannot enter a default time without entering a default date at the same time, but you can enter a default date without an accompanying default time. Use a hyphen (-) or slash (/) as date separators; for example 10/24/2000 or 10-24-2000. Use a colon (:) as a time separator; for example, 11:59:00.

   The default date and time value automatically becomes the value for the date/time field for all existing or newly created items for which this field is a property. You can change the value manually for any item on the Custom tab on the *Property* dialog.

   If you do not set a default, no value is placed in the database for any item with this field until you set the field to a value manually.
3. Click OK.

Creating a User ID Type

This procedure is a continuation of "Creating a User-Defined Property Field" on page 180. Remember that you cannot set a default value for a user ID field because if a user ID is deleted the selected value will become invalid.

To create a user ID type:

1. Select User ID from the Type text box.
2. Select the Input Required check box if you want to make this a required field.

Note Required fields are meaningless when you add, check in, and check out files. Because they affect only the File Properties dialog, required fields are rarely used for files.

3. Click OK.

Modifying an Existing Field

You can modify some existing property fields and any custom fields that have been created. The property fields that can be modified are always enumerated types.

For example, Priority is an existing change request property that can be customized. On the Synopsis tab of the *Change Request Properties* dialog, Priority has the values Yes and No, but is implemented as an enumerated type with Yes having the numeric code of 1 and No having the numeric code of 0. If your company prefers to prioritize change requests on a scale from 1 (high) to 10 (low), you can add additional values. You can use the current numeric codes 0 and 1 with new names or disable them.

Try to take advantage of the existing codes in some way.

For example, you might change the name of the enumerated value No to Not Prioritized. Then you can query for Priority equal to 0 to identify the change requests that need to be prioritized.

You might also change the name Yes to Priority 1, as it already has the code 1. Then you can add additional values for Priority 2 through Priority 10. Any numeric codes that you assign must have a value greater than 100, because the application reserves the values 0 through 100 for internal uses. So you might add a value with the name Priority 2 and the numeric code 101, Priority 3 with the numeric code 102, and so on.

Queries use the numeric codes to make comparisons based on enumerated values. Using the above codes, you can write a query to locate Priorities 1, 2, and 3, by querying for Priority greater than 0 and less than 103.

By default, sort and list boxes display items in the order that they appear in the *Add Field* or *Modify Field* dialog. You can use the sort options to sort by name in ASCII or alphabetical order.

To modify an existing field:

1. Select Advanced > Customize from one of the item menus.

   The Customize dialog lists application fields that can be customized by an icon with a pencil in the center it (fully customizable) or an icon with a small yellow lock in the bottom left corner (where only display names can be changed). These fields are considered enumerated and restricted types, respectively.

   If user-defined fields are enabled, it shows an icon with the head and shoulders of a person, and if they are disabled, the same icon is greyed-out. All user-defined fields are fully customizable.

Note If a field is disabled, it appears in the *Show Fields* dialog only when the Show Advanced Fields check box is selected.

2 Select the field to be edited. Click Edit. The *Modify Field* dialog appears.

Tip Double-clicking the name of the field also opens the *Modify Field* dialog.

3 Follow the steps in Table 11.2, “Modifying Existing Application Fields” as you edit. Each type of field has different editable characteristics.

4 After making your changes, click OK.

Table 11.2 Modifying Existing Application Fields

Permitted Modifications	Modifications Not Supported
Enumerated Type	
To add values to possible values, see “Creating an Enumerated Type” on page 181 for more information.	Cannot change the field name.
To change the names of possible values, select a value, click Edit, and type a new name in the Name text box.	Cannot change the display name.
To change the default value, select a new value from the Default Value list box	Cannot change the type.
To enable disabled values, select a disabled value and click Enable.	Cannot change the numerical codes associated with the possible values.
To reorder possible values in the enumerated list in the *Add Field* and *Modify Field* dialogs, use and drop.	Cannot change the numerical codes associated with the possible values.
Restricted Enumerated Types	
To change the names of possible values, select a value, click Edit, and type a new name in the Name text box.	Cannot change the field name, the display name, or the default value. Cannot add, disable, or reorder possible values.

Table 11.3 Modifying User-Defined Fields

Permitted Modifications	Modifications Not Supported
Enumerated Type	
To change the display name, enter a new name in the Display Name text box.	Cannot change the field name.
To add new values to possible values, see "Creating an Enumerated Type" on page 181 for more information.	Cannot change the type.
To change the names of possible values, select a value, click Edit, and enter a new name in the Name text box..	Cannot change the numerical codes associated with the possible values.
To disable possible values, select an enabled value and click Disable.	Cannot change the numerical codes associated with the possible values.
When a field is disabled, it appears only in the *Show Fields* dialog when the Show Advanced Fields check box is selected.	
To enable disabled values, select an enabled value and click Disable.	Cannot change the numerical codes associated with the possible values.
To change the default value, select a new value from the Default Value list box.	
To reorder the possible values in the enumerated list in the *Add Field* and *Modify Field* dialogs, use drag and drop.	
Integer or Real Types	
To change the display name, enter a new name in the Display Name text box.	
To change the default value, enter a new value in the Default Value text box.	

Table 11.3 Modifying User-Defined Fields (continued)

Permitted Modifications	Modifications Not Supported
Text Type	
To change the display name, enter a new name in the Display Name text box.	Cannot change the maximum length.
To indicate whether the field is required, select or clear the Input Required check box.	
Date/Time Type	
To change the display name, enter a new name in the Display Name text box.	
To change the default value, enter a new value in the Default Value text box.	
User ID Type	
To change the display name, enter a new name in the Display Name text box.	
To indicate whether the field is required, select or clear the Input Required check box.	

Translating Display Names and Enumerated Values

When you create a custom field, you provide a display name for that field. If you create the field is an enumerated type, you must also provide names for the enumerated values. All of these names are also displayed to users in the client.

The names that you specify when you create the field become the default names for that field.

If you do not set up translations for a custom field's names, users will always see the default names. However, if you translate the names associated with a field, users will see the names provided for the locale specified for the client's operating system.

For example, suppose you have a custom field for which all the default names are in English. Then you translate the names (a display name and several names for enumerated values) into French. Users whose clients use French operating systems will see the French names. All other users see the default names.

Note Although you can set up translations from the Windows and Cross-Platform clients, only the Cross-Platform clients has been localized. Therefore, your translations will be seen in the client interface only by users of those clients. Users of the Windows client will always see the default names, except during the process of translating names for custom fields.

In more complicated examples, you need to understand how locales and languages are defined. Each locale or language is defined by using at most three parts separated by hyphens. These parts represent the language itself, the country it is spoken in, and a variant of that language. Most languages are represented by just the first two of these parts.

For example, Dutch has the following three entries in the language table:

Dutch [nl]

Dutch (Belgium) [nl-BE]

Dutch (Netherlands) [nl-NL]

The bracketed section represents the locale or language. In this case, [nl] represents a generic Dutch, while [nl-BE] and [nl-NL] represent Dutch as it is spoken and written in Belgium and in the Netherlands, respectively.

If you translate the names from a custom field into generic Dutch, the names can be used for all Dutch operating systems. If you translate the names into both generic Dutch and the Dutch used in Belgium, here's what happens:

- Users whose clients run on operating systems with the locale set to the Dutch of the Netherlands (the [nl-NL] entry) see the generic Dutch translations.

  If an enumerated name is not translated into Dutch, these users see the default name for that value.

- Users whose clients run on operating systems with the locale set to the Dutch of Belgium see the translations for the Dutch of Belgium.

  If an enumerated name is not translated into the Dutch of Belgium, they see the generic Dutch name for that value. If a name is not translated into either the Dutch of Belgium or generic Dutch, they see the default name for that value.

  Thus, if you are generally satisfied with the generic Dutch translations, but want to use a different spelling or a different word for just one value in the Dutch of Belgium, you can translate the name of that value into the Dutch of Belgium.

Note The Java VM reads the locale information setting for the client's operating system when it starts running. If you change that setting, you must to restart the Java VM.

Important Generally, an organization standardizes on one language for the default values of all custom fields. Otherwise, the application allows you to use a different default language for every custom field. For example, when you create a custom field using English names, English names become the default for that field. If someone else creates a custom field using Japanese names, the default names for that field are Japanese, unless they are changed later.

To translate the display name and/or enumerated values for a custom field:

1. Do one of the following:
   - While creating a custom field with the *Add Field* dialog, click Translate.
   - Display the properties of a custom field:
     1. Select the Advanced > Customize command from one of the item menus, such as the File or Change Request menu.
     2. Select the field to be modified.
     3. Click Edit, to display the default values in the *Modify Field* dialog. These values may not be what you see in your client. For example, the defaults may be in English, and you may be using the Cross-Platform client with a French locale.
     4. Click Translate. The *Translate Field* dialog appears.
2. If the language you want appears in the Language drop-down list box, select it.

   If it does not appear, click Add.

   a. From the *Add Language* dialog:
      1. Select a language from the Language list box.
      2. In the User Defined text box:
         - Enter an existing locale name, such as fr-FR to quickly select the language by its locale name.
         - Enter a new locale name. A user-defined locale name must follow the formatting rules for locale names. You can use up to 8 characters: letters, numbers, and one or two hyphens to separate the locale name into two or three parts. Underscores display as hypnens. (No spaces are allowed.) Regardless of the case you enter, the first part will always be lowercase and the second part will always be uppercase. After its creation, the user-defined language becomes a member of the language list for the current

field. It is not available for other fields, unless you also create it for those fields. To see the values in a user-defined language, users must set their workstations to that locale.

b. Click OK, to return to the *Edit Translations* or *Translate Field* dialog.

Be aware that adding a language to one field does not automatically add it to the language list for any other field. Also, if you do not translate any values, the language will disappear from the language list for this field.

3. (Optional) In the Translated Display Name text box, enter the display name for the field in the new language.
4. (Optional) Translate enumerated values:
   a. Select an enumerated value to be translated.
   b. Click Edit. The *Translate Enumerated Value* dialog displays.
   c. Enter the translated value in the Translated Text box.
   d. Click OK to return to the *Translate Field* dialog.
   e. Repeat steps a through d for other enumerated values.
5. Click OK to return to the *Add Field* or *Modify Field* dialog.
6. Click OK to return to the *Customize* dialog.
7. Click Close.

Chapter 12

Controlling Access to Objects

By default, all users initially have access to everything in the application. To avoid accidental deletions and other problems, administrators must set access rights as soon as possible.

Access rights can be categorized as:

- Server-level access rights

  These access rights allow users to perform server administration operations, such as modifying server configurations and viewing logs.

  Additional rights at the server level include the rights to create projects, create custom fields, control component-level access rights, and perform certain operations specific to the Notification Agent.

- Project, view, folder, and item-level access rights

  Initially, any user who can see one of these objects can set rights for it. However, project-level, view-level, folder-level, and even item-level rights function hierarchically and may be affected by group privileges and ownership. See “Understanding Access Rights for Projects, Views, Folders, and Items” on page 192 and “Understanding Access Right Levels” on page 195 for more details.

- Component, filter, and query-level access rights

  The application components (file, change request, requirement, task, and topic) are server-wide objects. For example, the change request component appears in every project view and has the same filters and queries in every view.

  Component-level access rights govern the use of filters and queries for each component. They determine the users who can create public filters and queries in that component, who can use the filters and queries, and so on. A server-level access right named “Administer component-level access rights” allows users to set these rights.

  Individual filters and queries also have access rights. These rights override the general access rights set for filters and queries. Also, if ownership is not ignored, the user who creates a public filter or query is considered the owner of the filter or query and has access to it.

  The right pane contains a tree of access rights subcategories. When expanded, each subcategory displays a set of access rights as its children.

Setting Server-level Access Rights

The server-level rights you assign to users and groups authorize them to perform specific operations in a particular server configuration. One of the options determines who can and who cannot create projects when the server configuration is running.

Server rights can be assigned only when a server is running.

By default, the Administrators group is assigned all project and Server rights. By default, the All Users group has the rights to create projects and review the server configuration and the server log. The Server access rights are briefly described in the following table.

Table 12.1 Server Access Rights

Server Operations	
This access right...	**Allows a user or group to...**
View server log	Review, but not change, server log information.
View statistics and licensing information	Review, but not change, statistics information (StarTeam Server 5.4 and earlier). Create license usage files.
View server configuration	Review, but not change, the server configuration options.
Modify server configuration	Change the server configuration options.
Remotely administer server	Lock/unlock the server, restart the server from the client, shut down the server from the client, access the Start/Stop Conversion and Hive Manager vault buttons.
Administer user accounts	Add groups and users.
View system policy	Review, but not change, the password and logon failure options for the server configuration.
Modify system policy	Change the password and logon failure options for the server configuration.
Change server security settings	Set Server access rights. If you change this setting, be sure that you remain one of the users who can change access rights.
View security log	Review, but not change, server log information
StarDisk Operations	
This access right...	**Allows a user or group to...**
Create new users	Add new users to sample project.
Replication Support	
This access right...	**Allows a user or group to...**
Change user/operation time	Manipulate creation times and user names when using special clients, such as Notification Agent.
Project Operations	
This access right...	**Allows a user or group to...**
Create projects	Create projects when the Server is running the server configuration.
Customizations	
This access right...	**Allows a user or group to...**
Add/modify database schema	Create customized fields as item properties, or modify a field for an item that can be modified.
Component operations	
This access right...	**Allows a user or group to...**
Administer component-level access rights	Designate the users and groups who can create and apply filters and queries for a specific component in the server configuration.

Table 12.1 Server Access Rights (continued)

Component Operations	
This access right...	**Allows a user or group to...**
Administer component level rights	Set access rights at the component level.

To set Server access rights:

1 Display the *Server Administration* dialog by doing one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).
- From the computer on which the Server is installed, enter the following at the command line: `serveradmin`

2 Select a server configuration from the left pane. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 Do one of the following:

- From the lower left pane, click the Accounts bar. Then select the Access Rights shortcut.
- Select Tools > Accounts > Access Rights from the menu.

These actions display the *Access Rights* dialog.

4 Click New. This action displays the *Add a User or Group* dialog.

5 Select the user or group to be assigned access rights.

6 Select the Grant option button; then click OK to return to the *Access Rights* dialog.

Caution Never select Deny unless you are creating an exception. See “Denying Access Rights” on page 200.

7 Select or clear the appropriate check boxes.

The right pane contains a tree of subcategories, each of which has access rights as children. Selecting or clearing the check box for a subcategory such as “Server Operations” selects or clears all access right check boxes for that subcategory.

When the subcategory check box contains a check mark, the access right check boxes for that subcategory are selected. When it is cleared, the access right check boxes for that subcategory are either all cleared or mixed (some are selected and some cleared).

8 Click OK.

Caution Clicking Delete removes the selected user or group from the User and Groups list. As a result, the user or group loses any previously set access rights to the server.

Reviewing the Security Log

The Server records security events, such as adding users and groups, in the security log. For more information, see “Reviewing the Security Event Log” on page 45.

Understanding Access Rights for Projects, Views, Folders, and Items

As users log onto a server configuration, they are identified by their user names and as members of the groups to which they belong. This information is stored as an access token for each user.

When users perform operations on objects (projects, views, folders, and items), the application examines these tokens and the access rights for the objects on which the users are performing the operations.

Determining Access Rights and Tokens

The Server checks access rights in layers. The right to access an object begins with the *System Policy* dialog, which can be reached from the Server Administration dialog or the Server Administration Tool. See "Using or Ignoring Ownership and Privileges" on page 193 for more information about how to use the Access Rights tab of the *System Policy* dialog in relation to ownership and privileges.

By default, the application gives special rights to the owner of an item or object. Generally speaking, the owner is the user who created the project, view, folder, or item. Unless ownership is being ignored, the right of ownership overrides and takes precedence over all rights configured elsewhere.

Similarly, unless privileges are being ignored, the privileges granted to groups also override and take precedence over the access rights configured elsewhere. Privileges are group properties that are set by using the Privileges tab of the *Group Properties* dialog. See "Setting Group Privileges" on page 193 for more information.

A user is granted the same privileges as the group to which he or she belongs. If the user belongs to two groups and one group is granted certain privileges and the other group is denied the same privileges, the user is granted the privileges because at least one group to which he or she belongs has those privileges.

After checking ownership and privileges, the application checks the access rights granted for specific objects. Settings on the *Access Rights* dialogs for projects, views, folders, and individual items grant or deny users or groups the ability to perform operations at those levels.

Important: If rights are granted to any user or group at a given level in an *Access Rights* dialog, those users who are not granted rights at that level are effectively denied the rights.

Ultimately, if a user can see an object and no deny records stop the user from performing an operation, the user can do anything that a grant record allows him or her to do, whether as an individual user or as a member of a group. The only exceptions involve issues of ownership and privileges.

To summarize, the application performs the following checks to determine whether a user can perform an operation:

1. If the user requesting the operation is the owner of the object being operated on and ownership is not being ignored, access is granted.

   Note that ownership, when not ignored, takes precedence over both privileges and access rights.

2. If the user belongs to a group that has a satisfactory privilege and privileges are not being ignored, access is granted.

   Note that privileges, when not ignored, take precedence over access rights wherever access rights are set. If users belong to a group that has the correct privileges, they can be granted access rights that are specifically denied to them in the application.

3. If the user or any group to which the user belongs has been granted satisfactory access rights for the object on which the operation will performed, access is

granted. If the object has access rights set, but none are satisfactory, the user is denied access.

4 If the object has *no* access rights set, the application checks the next higher level.

For example, if the operation is on a file, change request, topic, task, or child folder, the application checks the access rights for the parent folder.

If the operation is on a root folder, the application checks the access rights for the view. If the operation is on a view, the application checks the access rights for the project. If the operation is creating a project, the Server access rights are checked.

5 If none of the levels has access rights set, access is granted.

Using or Ignoring Ownership and Privileges

As an administrator, you can override object ownership and/or group privileges by setting options from the server configuration *System Policy* dialog. Use these options with caution, because they change the steps used by the Server to check every user (including administrators) for access to all objects in the repository.

If you ignore both ownership and privileges, only access rights determine who can and cannot perform operations on objects in the repository.

To use or ignore object ownership and/or group privileges:

1 Display the *Server Administration* dialog by doing one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).
- From the computer on which the Server is installed, enter the following at the command line:

```
serveradmin
```

2 Select a server configuration from the left pane. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.

3 Do one of the following:

- From the lower left pane, select the Accounts bar. Then click the System Policy shortcut.
- From the menu, select Tools > Accounts > System Policy.

These actions display the *System Policy* dialog.

4 Select the Access Rights tab.

5 Set the desired combination of Access Validation Options:

- Select or clear the Ignore Object Ownership check box. When cleared, the server configuration checks for ownership.
- Select or clear the Ignore Group Privileges check box. When cleared, the server configuration checks for privileges.

6 Click OK.

Setting Group Privileges

The privileges assigned to a group may allow members of that group to access objects and perform operations that they are otherwise not allowed to do. In other words,

privileges override the access rights settings. See “Using or Ignoring Ownership and Privileges” on page 193 for information on how to stop the server configuration from checking for privileges.

If you select User Manager from the *Server Administration* dialog, you will notice that the server configuration comes with some default groups: All Users, Administrators, System Managers, and Security Administrators. The default user named Administrator belongs to both the Administrators and the Security Administrators groups.

By default, the Administrators group has all group privileges. Also by default, e other groups have none of these privileges.

All members of a group have the same privileges on every project managed by this server configuration. The privileges apply to all levels equally: projects, views, folders, and items within folders. If users belong to more than one group, they have the maximum amount of privileges, regardless of which group provides them with those privileges.

Table 12.2 Group Privileges

This privilege...	Allows a group to...
See object and its properties	See all projects, views, folders, items, and their properties. This privilege overrides the similarly named access right found in the Generic Object Rights in the Access Rights dialogs.
Modify object properties	Modify the properties of any projects, views, folders, or items. This privilege overrides the similarly named access right found in the Generic Object Rights in the Access Rights dialogs.
Delete object	Delete any projects, views, folders, or items. This privilege overrides the similarly named access right found in the Generic Object Rights in the Access Rights dialogs.
Purge object (delete permanently)	This privilege is not supported at this time.
Acquire object ownership	This privilege is not supported at this time.
Change object access right	Change access rights for any projects, views, folders, or items. This privilege overrides the similarly named access right found in the Generic Object Rights in the Access Rights dialogs.
Create object and place it in a container	Create new objects and put them in containers. When this privilege is set, the group can add new views to a project, new folders to a view, and new folders and items to a folder. This privilege overrides the similarly named access right found in the Generic Object Rights in the *Access Rights* dialogs. It does not override the server-level access right that allows users to create projects.
Grant all specific class-level rights for all classes of objects	Perform any operation not covered by the preceding privileges. For example, this privilege allows group members to check out files, break locks, perform linking operations, and perform labeling operations. This privilege overrides some of the access rights found in the Generic Object Rights and the Generic Object Container Rights and all of the access rights in the <item>-specific Rights in the Access Rights dialogs.

To set privileges:

1 Display the *Server Administration* dialog by doing one of the following:

- From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server.
- From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).
- From the computer on which the Server is installed, enter the following at the command line.

```
serveradmin
```

2. Select a server configuration from the left pane. If you have not yet logged on, you will be asked to do so. If you are using the client, you can administer remote servers only.
3. Do one of the following:
   - Click the Accounts bar in the lower left pane, and select the User Manager shortcut.
   - Select Tools > Accounts > User Manager from the menu.

   These actions display the *User Manager* dialog.
4. Add or select a group in the *User Manager* dialog. (See "Adding Groups" on page 52 for information.)
5. Add users to the group, if appropriate. (See "Adding Users" on page 55 for information.)
6. Right-click the group's name in the Groups tree and choose Properties from the context menu. This action displays the *Group Properties* dialog.
7. Select the Privileges tab.
8. Select or clear the check boxes to grant privileges to the group or take them away.
9. Click OK.

Understanding Access Right Levels

Access rights are defined for individual users or groups at the following levels:

- Project level: Access rights can be defined for the project itself. You can also define access rights that apply to all its views, child folders and items, unless a object has access rights set specifically for it. There are View, Child Folders, and other nodes at this and other levels.
- View level: You can define access rights for the view itself. You can also define access rights that apply to all its child folders and items, unless a specific object has access rights set specifically for it.
- Folder level: You can define access rights for the folder itself. You can also define access rights that apply to all its child folders and items, unless a specific object has access rights set specifically for it.
- Item level: You can define access rights to a specific file, change request, requirement, task, or topic. (It is unusual to set rights at this level.)

Note that project access rights can be set only at the project level, because that is the only level with Project node in the access rights hierarchy. You can set view access rights at either the project or the view level, because both of those levels have a View node. You can set folder access right at the project, view, or folder levels, and so on.

Opening Projects and Views

A project is indistinguishable from its initial view and also from the root folder of that view. In fact, any view of a project is indistinguishable from its root folder. Therefore, a user will not be able to open a project if you deny that user (or all groups to which the user belongs) any of the following:

- Ability to see the project.
- Ability to see the initial project view.
- Ability to see the root folder of the project's initial view

A user will not be able to open a particular view of a project if you deny that user (or all the groups to which the user belongs) any of the following:

- Ability to see that view.
- Ability to see that view's root folder.

Granting Access Rights at the Project Level

The following section provide information about setting access rights at the project level. It illustrates this information by using an example with three user-defined groups: Developers, Testers, and Others. (These groups are in addition to the All Users, Administrators, System Managers, and Security Administrators groups that come with the Server.) The example also assumes that the All Users group is larger that the Others group.

The Project Node

Assume that you decide that only members of the Administrators group should control the project and create a grant record. This record prevents anyone who is not a member of the Administrators group from seeing the project, unless ownership and privileges apply. As a result, no one else can access and work with the objects in this project.

Note Although members of the Administrators group require all access rights for the project, you may decide to omit them from the Users and Groups list if they have group privileges. Normally, this is acceptable. However, if the server configuration is set to ignore privileges, you must specifically grant the Administrators group all project access rights.

Next, you must assign the correct rights to the every other group that needs to access this project.

Because keyword expansion is a project property, the Developers group needs to have the rights to see the project and modify its properties. However, they probably do not need to delete the project or change its access rights.

The Testers and Others groups need to see the project and its properties, so you should select only the See Object And Its Properties check box for these groups.

The View Node

View access rights at the project level apply to all views that now exist or will be created for this project in the future. Members of the Administrators group need all view rights. They may be assigned these rights or receive them because of their privileges.

The Developers and Testers groups need to see and modify view properties and perform operations on labels. They do not need to create or delete views, manage promotion states, or change view access rights. The Others group needs to see the view, but requires no other rights:

The Promotion State Node

The Promotion State node is not important in this example. For more information, see "Understanding Access Rights for Promotion States" on page 212.

The Child Folders Node

For access rights to child folders at the project level, the Administrators group may need all rights. They may be assigned these rights or receive them because of their privileges.

The Developers and Testers groups probably do not need to delete folders, share or move folders, change folder behaviors or configurations, or change folder access rights.

You may want the Others group to only see the folders, their properties, and their histories.

The Item Nodes

Borland does not recommend creating only one grant or deny record for a given node. The following section illustrates how project-level item access rights work, using files as an example.

If only the developers need to access files, you can grant only the Developers group all file access rights at the project level.

With this setting, only members of the Developers group have access rights to any files, regardless of the view, folder, or file. As a result, only developers can see or perform operations on any files. Members of the Testers and Others groups see only the files that they have in working folders, but the status of these files appears as Not In View. However, by default, two exceptions exist:

- Ownership. Suppose a user named MBenson is a tester but also is the owner of a file. That means that MBenson's name is in the Created By field for that file. By default, MBenson always has complete access to that file.
- Privileges. If groups other than the Developers have one or more privileges that allow them to see, modify, define, or perform other actions on a file, members of those groups have access to the files regardless of the access rights settings. For example, the default settings for the Administrators group grant all privileges to this group. Therefore, members of this group can perform any file operations.

You can stop the server configuration from checking for either ownership or privileges or both. See "Using or Ignoring Ownership and Privileges" on page 193 for more information.

If you give only the Testers and Developers groups access to other types of items (change requests, requirements, tasks, and topics), the same exceptions apply.

However, other groups will want to see these types of items, so you will need to grant these groups some access rights.

Granting Access Rights at the View Level

Usually, granting access rights at the project level is not a fine enough level of granularity. For example, one set of developers may maintain Release 1.0 of the product in one view, while another set of developers writes the source code for Release 2.0 in another view.

To handle this situation, you may want to create new groups, such as 1.0 Developers, 2.0 Developers, 1.0 Testers, and 2.0 Testers. You can give the 1.0 Developers and 1.0 Testers access to files and/or change requests in the Release 1.0 view and. Then you can give the 2.0 Developers and 2.0 Testers access to files and/or change requests in the Release 2.0 view.

In this case, you need to set access rights at the view level. However, you must still set project access rights at the project level because that is the only place where the Project node appears.

View and Child View Access Rights

Access rights in a child view *at the view level* are independent of the access rights of the parent view *at the view level*. Therefore, a child view starts out with no access rights at the view level.

A new child view is represented by a different object in the repository from the parent view. It has a different name, description, place in the view hierarchy, etc.

View-level access rights can be set for a new child view. For example, suppose a reference view contains only one branch of the parent view's folder hierarchy. The reference view has a root folder named QA Tests. In this situation, you can make the Testers the only group with file access rights in the reference view, even if Developers is the only group that has file access rights in the parent view.

Access Rights at Different Levels

Sometimes a group has different access rights at the view and the project levels for the same type of object in the same view. In this situation, the access rights at the lowest level are enforced.

When the Server searches for access rights, it starts from the lowest level and moves to the highest level. When it finds a level at which a group has access rights, it does *not* search any higher levels for that type of object.

Remember that the project access rights exist only at the project level, so the project level is always searched for these rights. File access rights, on the other hand, exist at the file, folder, view, and project levels. the Server stops at the first level at which it finds file access rights.

Granting Access Rights at the Folder Level

Setting access rights at the folder level is usually done when you want to allow certain groups (but not other groups) to access a particular branch of the folder hierarchy. For example, you may want only the Writers group to be able to access the branch that has User Manual as its root folder.

Setting access rights at the folder and the item levels has more consequences than setting rights at higher levels. When a child view is derived from a parent view, as all reference and most branching views are, it initially contains objects that belong to its parent. In branching views, these objects can branch into new objects that exist only in the child view. Just as a new view has no view-level access rights, folders and items that branch into new objects initially have no access rights at the folder or item level.

This Folder and Child Folder Nodes

The folder level has two nodes—This Folder, for the selected folder, and Child Folders, for the other folders in the folder hierarchy of the branch. This feature allows you to set different access rights for each node.

The Root Folder

In the application, the root folder of a view can be indistinguishable from that view. If the view is the root (or initial) view of a project, the root folder can be indistinguishable from that project. For more details, see "Managing Folders" in the *StarTeam User's Guide*.

Using the This Folder node to set access rights for the root folder can therefore affect a user's access to a view. If the view is the root view, it can also affect the user's access to the project. Therefore, most administrators avoid setting folder-level access rights on a root folder, as these rights may interfere with view-level or project-level rights.

For example, suppose the Developers group is not granted the right to see the User Manual folder and that this folder is the root of a reference view. Then members of the Developers group cannot open that view, even if view-level access rights allow them to see the view. An error message appears when they try to open the view. Users who can see a project but not its root view also see an error message.

Access Rights of Child Views

If a child view includes child folders that have access rights in the parent view, its access rights depend upon whether it is a reference view or a branching view.

The access rights in a reference view at the folder level are *not* independent of the access rights at the folder level in the parent view, as no branching will ever occur. You can see these access rights from either view, if you have the rights to do so.

If you change access rights in the reference view, you simultaneously change the access rights in the parent view (and vice versa) because the folder in the reference view is the same object as the folder in the parent view.

If the child view is a branching view, the access rights in the child view at the folder level are independent of the access rights at the folder level in the parent view, but only after the folder in the branching view actually branches.

Initially, any folder you see in the branching view is the same object that exists in the parent view. Therefore, it has the same access rights in both views. Initially, you can change access rights in the parent view (and vice versa), because the folder in the branching view is the same object as the folder in the parent view. Once the folder branches, however, a new object for that folder is created in the branching view. This object begins a life cycle of its own and no longer has any access rights at the folder level.

Note Remember that branching a folder does not branch any of the folder's contents. Each item is the folder is treated separately.

Folders in a branching view have something that folders in a reference view do not have: behavior.

If a folder branches on change and you change one of its properties, its revision number changes. When the folder branches, it becomes a new object in the repository and no longer has any access rights at the folder level.

If a folder does not branch on change and you change one of its properties, the revision number changes, but no new object is created. In this case, the folder retains its access rights in both views. Because both views still contain the same object, changes you make to the folder's access rights in one view also change that folder's access rights in the other view.

Granting Access Rights at the Item Level

Although access rights can be set on individual items, this is rarely done. For example, if you really need to allow only one person to know about a particular file, you can give only that person access rights to that file. However, by default, the owner of the file and anyone belonging to a group with the correct privileges can still see that file.

To ensure that only that one person can access the file, you would have to stop the Server from checking for privileges and ownership. Then the access to every object would be controlled solely by access rights.

Like folders, items in a child view retain the access rights they had in the parent view until they branch into new objects. Items lose their access rights only when branching.

Moving Folders or Items

When you move a folder or an item, the access rights set for it at the folder or item level go with it. For example, if you move the User Manual folder in the StarDraw view to another view, it has the same folder-level access rights in the new view that it had in the StarDraw view. It also has the same behavior, which either allows it or stops it from branching on change.

Sharing Folders or Items

When you share a folder or item, the access rights set for it at the folder or item level accompany it, until the folder or item branches.

When you share a folder or item, its behavior may change. When shared, the behavior immediately becomes able to branch on change, even if the Branch On Change check box was disabled in the original location. Whether the Branch On Change check box is selected or cleared depends on the property setting for the destination view Set Items That Are Shared Into View To Branch On Change.

When the folder or item branches in its new location, a new object is created in the repository, and that new object initially has no access rights at the folder or item level.

Automatically Granting Access Rights

The Server provides a certain degree of automatic assistance with access rights. For example, if you allow a group to check in files, but do not grant its members the right to see the files or file histories, check out the files, change file properties, or lock the files, the Server automatically gives the group those rights. The system grants these rights because members of the group could not check in files without them. Although you exited the *Access Rights* dialog without selecting those check boxes, you will see checkmarks in them the next time you open the dialog and look at the access rights node.

Denying Access Rights

Deny records are rarely used. However, they do allow you to create exceptions to the current access rights.

Creating a grant record with no check boxes selected is not the same as creating a deny record with all the check boxes selected, although both stop users and groups from performing the same operations.

For a given node at a given level, grant records are examined until one gives a user or group permission to perform an operation or until all the grant records have been examined without finding one that gives permission. If membership in one group does not allow a user to perform an operation but membership in a second group does, the user can perform the operation.

However, if a deny record for that node forbids the user from performing an operation, the user cannot perform that operation. The application disregards any grant records for the same node that allowed the user to perform the operation.

Important Be aware that the Windows client and the Cross-Platform client interpret the Deny Access Rights checkboxes differently. Unchecked items are interpreted as valid in the Windows client, while checked items are interpreted as valid in the Cross-Platform client.

Deny records must precede grant records. The reason is that if the application finds a grant record that allows a user to perform an operation before it finds a deny record for the user, it stops looking at records for that node at that level. Thus, it allows the user to perform the operation.

Note Object ownership or group privileges can override either grant or deny records.

General Access Rights Rules

- Access rights can be overridden by:
  - The fact that a user is the object's owner. Usually, the owner is the person who created the object.

  - Privileges given to a group that includes the user. These privileges are set per group from the Server. By default, the Administrators group has full privileges (rights to do anything and everything).
- Access rights should be set at the highest possible level.
- The application checks for access rights from the lowest level (the item level) to the highest level (the project level).
- If one grant record is created for a node, a grant record for that node should be created for every group that requires access to the project at that level. The Administrators group should have a grant record for each node, so that, if privileges are ignored, administrators can still change access rights.
- If access rights are set for any user or group for a node, all users or groups without a grant record for that node will be denied all access rights at that level for that node.
- Every view within a project has the same project-level access rights.
- When you derive a child view from an existing view, the new view has no view-level access rights. However, folders and items in the child view that existed in the parent view retain the same folder-level or item-level access rights that they had in the parent view.

  Changing these access rights in either the parent or the child view also changes them in the other view because you are changing the rights on the same object. If the folders or items in either the parent view or the child view branch, they can have different access rights, because they are different objects.
- Folders that are moved or shared from one view to another retain any access rights assigned to them at the folder level in the new view. However, if they branch, they lose their folder-level access rights.
- Items that are moved or shared from one view to another retain any access rights assigned to them at the item level to the new view. However, if they branch, they lose their item-level access rights.
- Avoid setting item-level access rights.
- Avoid creating deny records. But if you deny rights, follow both of these rules:
  - Never allow any node on an *Access Rights* dialog to have only deny rights. records.
  - Verify that deny rights records for a node precede any grant rights records for the node.

Setting Access Rights

The following procedure is a set of generic instructions for setting access rights for projects, views, folders, and items. The next few sections display the dialogs for each set of access rights and explain them.

To set access rights:

1 Select the object or item for which rights will be set.
  - To set project access rights, the project must be open to any view.
  - To set view access rights, the view must be open.
  - To see folder access rights, select the folder from the folder hierarchy in the left pane.
  - To set file, change request, requirement, topic, or task access rights, select the item from the upper pane.

2 Select Access Rights from the appropriate menu or context menu.

- To set project-level access rights, select Project > Access Rights from the menu bar.
- To set view-level access rights, select View > Access Rights from the menu bar.
- To set folder-level access rights, select Folder > Advanced > Access Rights from the menu bar.
- To set item-level access rights, select <item> > Advanced > Access Rights from the menu bar.

3 From the *Access Rights* dialog, click Add to select a user or group. The *Assign Access Rights To* dialog appears.

4 Select a user or group. Users are listed by their user names and groups are listed by their paths (excluding the All Users group).

5 Select the Grant option button.

6 Click OK to return to the *Access Rights* dialog.

Caution Never select the Deny option button unless you are creating an exception. See "Denying Access Rights" on page 200. Deny records must be created before Grant records.

7 Select and/or clear the appropriate check boxes. Selecting or clearing the check box for a category, such as "Server Operations" selects or clears all the access right check boxes for that category.

The category check box has only two states. When it is cleared, the access right check boxes for that category are either all cleared or mixed: some selected and some cleared.

8 Click OK.

Caution Clicking Delete removes the selected user or group from the User and Groups list. The selected user or group loses any previously set access rights to the server.

If you grant a user or group the ability to modify an item, the user or group must also have the ability to view the item. The application assists you with this task. For example, if you select only the Modify File Properties check box and click OK, when you reenter the dialog, the application will also have selected the See File And Its Properties check box.

Setting Project Access Rights

The following table describes the generic object rights for a project. To display the *Project Access Rights* dialog, select the Project > Access Rights command.

The right to create a project is set as a Server access right. See "Setting Server-level Access Rights" on page 190.

Table 12.3 Project Access Rights: Generic Object Rights

This access right	Allows a user or group to...
See object and its properties	See this project and view its properties by selecting Project > Properties.
Modify properties	Change the properties for this project. The project properties that can be modified are name, description, keyword expansions settings, alternate property editor (APE) settings, process rules settings, requiring unlocked files to be read-only, and several settings that affect users (for example, requiring revision comments to be entered when a file is checked in).
Delete object	Delete this project from its server configuration.
Change object access rights	Change the access rights for this project. If you change this setting, be sure that you remain one of the users who can change access rights.

Setting View Access Rights

When you select the View > Access Rights command to open the *View Access Rights* dialog, the rights shown are for the current view. The rights available from the View node are also available from the View node in the *Project Access Rights* dialog. In the latter case, the rights cover all views in the project rather than an individual view. It also include a container-level right that allows users or groups to create views for the project. This right is not available on the View node of the View Access Rights dialog.

The following table describes the access rights that are available from the View node in the *Project Access Rights* dialog. Most of these access rights also appear on the View node of the *View Access Rights* dialog, but apply only to the current view.

Table 12.4 View Access Rights

Generic Object Rights	
This access right...	**Allows a user or group to...**
See object and its properties	See views in the views list (upper pane) and view its properties by selecting View > Properties.
See object and its properties	Change view properties. View properties that can be modified are the view's name, description, working folder (also the root folder's working folder), branch setting for shared items, and file status repository setting.
View-Specific Rights	
This access right...	**Allows a user or group to...**
Create view labels	Create view labels. These labels will be automatically attached to the folders and items in the view. Users with this right but not the right to attach labels can still create labels.
Modify view labels	Change the properties of view labels. For example, this right allows a user to freeze labels so that they cannot be adjusted
Delete view labels	Delete view labels. This action automatically detaches the view labels from the folders and items that had the labels. Users with this right but not the right to detach view labels can still delete view labels.
Create revision labels	Create revision labels. Users with this right but not the right to attach labels can still create labels.
Modify revision labels	Change the properties of revision labels. For example, this right allows a user to freeze labels so that they cannot be adjusted.
Delete revision labels	Delete revision labels. This action automatically detaches the labels from the folders and items that had those labels. Users with this right but not the right to detach revision labels can still delete revision labels.
Define promotion model	Create, delete, and reorder promotion states and edit their properties. After creating a promotion state, you must exit and reenter the *Promotion* dialog if you want to set access rights for the newly created state.
Create views	Create views in the current project. This container-level right is available only when you select the View node from the *Project Access Rights* dialog.

Setting "This Folder" Access Rights

When you select the Folder > Advanced > Access Rights command to display the *Folder Access Rights* dialog, you see two folder nodes. The rights available from This Folder node apply to the selected folder only. The rights available from the Child Folders node apply to all the child folders of the selected folder. See "Setting "Child Folders" Access Rights" on page 206. The dialog and following table refer to the current folder. The table describes the access rights that are available from the This Folder node in the *Folder Access Rights* dialog.

Note Because This Folder has no Generic Item Container subcategory for access rights, container rights for This Folder are on its Child Folders node. If This Folder is the root folder, these rights are set on the Child Folders node of the *View Access Rights* dialog.

Table 12.5 Folder Access Rights

Generic Object Rights	
This access right...	**Allows a user or group to...**
See item and its properties	View this folder's Name, Exclude, and Files tabs, which become available when Folder > Properties is selected. The History tab is controlled by the "See folder history" access right. The Link tab is controlled by the "See folder links" access right.
Modify properties	Change folder properties on the folder Name and Exclude tabs. Properties include folder name, description, use of inherited and local exclude lists, and contents of the local exclude list. If the folder is not a root folder, the working folder and alternate working folder settings are also properties. For root folders, the working folders are view properties and not controlled by this access right.
Delete from folder	Delete this folder from its parent folder. Be aware that if you can delete any of this folder's parent folders, you can still delete this folder.
Change item access rights	Change the access rights for this folder. If you change this setting, be sure that you remain one of the users who can change access rights.
See history	See this folder's History tab, which is available when Folder > Properties is selected.
Label Rights	
This access right...	**Allows a user or group to...**
Attach/Adjust view labels	Add a view label to this folder. Move a view label from one revision of this folder to another. This right controls direct manipulation of labels for this folder at the folder level. It does not stop users from attaching a view label to this folder when a view label is created.
Detach view labels	Remove a view label from this folder. Be aware that if users can delete view labels, they can detach a view label from this folder by deleting the view label from the view, regardless of the setting for this right.
Attach/Adjust revision labels	Add a revision label to this folder. Move a revision label from one revision of this folder to another. This right controls direct manipulation of revision labels for this folder at the folder level.
Detach revision labels	Remove a revision label from this folder. Be aware that if users can delete revision labels, they can detach a revision label from this folder by deleting the revision label from the view, regardless of the setting for this right.
Link Rights	
This access right...	**Allows a user or group to...**
See links	See the links involving this folder.
Create links	Link this folder to other folders and items.
Modify links	Change a link for this folder.
Delete links	Delete a link for this folder.

Setting “Child Folders” Access Rights

When you select the Child Folders node from the *Folder Access Rights* dialog, the available rights apply to the child folders of the selected folder. The Child Folders node is also available from the *View Access Rights* dialog and the *Project Access Rights* dialog. In these cases, the rights apply to all child folders in the current view or all the child folders in the project, respectively.

The next table describes the access rights available from the Child Folders nodes in the *Project Access Rights*, *View Access Rights*, or *Folder Access Rights* dialogs.

Table 12.6 Child Folder Access Rights

Generic Item Rights	
This access right...	**Allows a user or group to...**
See item and its properties	See the selected folder’s child folders or the selected project’s or view’s folders in the folder hierarchy in the left pane on the screen. You can also view the *Name* and *Exclude Properties* dialogs, which open when Folder > Properties is selected. The History tab is controlled by the “See folder history” access right.
Modify properties	Change folder properties on the Name and Exclude tabs for child folders. The properties include the folder’s name, description, use of inherited and local exclude lists, and the contents of the local exclude list. If a child folder is not a root folder, the working folder and alternate working folder settings are folder properties. If it is the root folder, the working folders are view properties and not controlled by this access right.
Delete from folder	Delete the selected folder’s child folders or the selected project’s or view’s folders from their parent folders. Be aware that if you can delete any of this folder’s parent folders, you can still delete this folder.
Change item access rights	Change the access rights for the selected folder’s child folders or the selected project’s or view’s folders. If you change this setting, be sure that you remain one of the users who can change access rights.
See history	See the History tab, which is available when Folder > Properties is selected. This action applies to the selected folder’s child folders or the selected project’s or view’s folders.

Table 12.6 Child Folder Access Rights (continued)

Label Rights	
This access right...	**Allows a user or group to...**
Attach/Adjust view labels	Add a view label to the selected folder's child folders or the selected project's or view's folders. Move a view label from one revision of a child folder to another. This right controls direct manipulation of view labels for child folders at the folder level. It does not stop users from attaching a view label to child folders when a view label is created.
Detach view labels	Remove a view label from the selected folder's child folders or the selected project's or view's folders. Be aware that if users can delete view labels, they can detach a view label from child folders by deleting the view label from the view, regardless of the setting of this right.
Attach/Adjust revision labels	Add a revision label to the selected folder's child folders or the selected project's or view's folders. Move a revision label from one revision of a child folder to another. This right controls direct manipulation of revision labels for child folders at the folder level.
Detach revision labels	Remove a revision label from the selected folder's child folders or the selected project's or view's folders. Be aware that if users can delete revision labels, they can detach a revision label from this folder by deleting the revision label from the view, regardless of the setting of this right.
Link Rights	
This access right...	**Allows a user or group to...**
See links	See the links involving the selected folder's child folders or the selected project's or view's folders.
Create links	Link the selected folder's child folders or the selected project's or view's folders to other folders and items.
Modify links	Change a link for the selected folder's child folders or the selected project's or view's folders.
Delete links	Delete a link for the selected folder's child folders or the selected project's or view's folders.
Generic Item Container Rights	
This access right...	**Allows a user or group to...**
Create and place in folder	Create a folder in a parent folder, view, or project in which the *Child Folder Access Rights* dialog has this option.
Share/Move out of folder	Share or move a folder in a parent folder, view, or project if its *Child Folder Access Rights* dialog has this option. Be aware that the access rights set for that folder and its contents, along with any rights set for specific child folders and items within that branch of the folder hierarchy, accompany the folder into the new folder.
Change behavior or configuration	Change the branching ability and configuration of folders that reside in a parent folder, view, or project if its *Child Folder Access Rights* dialog has this option.

Setting File Access Rights

When you open the *File Access Rights* dialog and select the File node, the rights apply only to the selected file.

File access rights are also available from the *Folder Access Rights* dialog, *View Access Rights* dialog, and *Project Access Rights* dialog. In these cases, the rights cover all files in the selected folder, view, or project rather than an individual file. Only in these cases are the following container-level access rights available:

- Add files to a folder
- Share/move file out of a folder
- Change file behavior/configuration

The following table describes the access rights that are available from the File nodes in the *Project Access Rights*, *View Access Rights*, *Folder Access Rights,* and *File Access Rights* dialogs..

Table 12.7 File Access Rights

Generic Item Rights	
This access right...	**Allows a user or group to...**
See item and its properties	See files in the files list (upper pane) and view file properties by selecting File > Properties.
Modify properties	Change the file properties. The properties that can be modified depend on whether the file is in Native-I or Native-II format. They include the archive /file name, description, executable bit setting (useful only for non-Windows platforms), compression, storage options, and custom properties. If used, an alternate property editor (APE) may restrict the properties that can be modified and the users who can modify them still further.
Delete from folder	Delete files from their folders.
Change item access rights	Change access rights for the files. If you change this setting, be sure that you remain one of the users who can change access rights.
See history	See file history in the history pane.
Perform maintenance	Change the revision comments for past revisions.
	Convert archives from Native-I to the Native-II format.
Set exclusive locks	Lock files exclusively.
Break exclusive locks	Remove someone else's exclusive lock on the files.
Label Rights	
This access right...	**Allows a user or group to...**
Attach/Adjust view labels	Add a view label to the files. Move a view label from one revision to another.
	This right controls direct manipulation of view labels for the files at the item level. This right does not stop users from attaching a view label to the files when a view label is created.
Detach view label	Remove a view label from the files.
	Be aware that if users can delete view labels, they can detach a view label from the files by deleting the view label from the view, regardless of the setting of this right.
Attach/Adjust revision labels	Add a revision label to the files. Move a revision label from one revision to another.
	This right controls direct manipulation of revision labels for the files at the item level. It can stop users from checking in files with an attached revision label.
Detach revision labels	Remove a revision label from the files. Be aware that if users can delete revision labels, they can detach a revision label from the files by deleting the revision label from the view, regardless of the setting of this right.
Link Rights	
This access right...	**Allows a user or group to...**
See links	See the links created for the files.
Create links	Link the files to other items.
Modify links	Change a link to the files.
Delete link	Delete a link that affects the files.

Table 12.7 File Access Rights (continued)

File-Specific Rights	
This access right...	**Allows a user or group to...**
Check in file	Check in the files.
Check out file	Check out the files.
Generic Item Container Rights	
This access right	**Allows a user or group to...**
Add file to folder	Add files to a folder, view, or project if its *File Access Rights* dialog has this option. This right appears only on the *File Access Rights* dialogs associated with a folder, view, or project.
Share/Move out of folder	Share files or move files that reside in a folder, view, or project if its *File Access Rights* dialog has this option. This right appears only on the *File Access Rights* dialogs associated with a folder, view, or project. Be aware that the access rights set for any file that is moved or shared go with that file into the new folder.
Change behavior or configuration	Change the branching ability and configuration of files that reside in a folder, view, or project if its *File Access Rights* dialog has this option. This right appears only on the *File Access Rights* dialogs associated with a folder, view, or project.

Setting Access Rights for Other Items

The access rights for change requests, requirements, tasks and topics are identical. If you select the <item> > Advanced > Access Rights command to open the *Access Rights* dialog, the rights apply only to the selected change request, requirement, task, or topic.

The < *item*> *Access Rights* dialogs are also available from the Folder > Advanced > Access Rights, View > Access Rights, and Project > Access Rights commands. In these cases, the rights cover all items of the specified type in the selected folder, view, or project. Only from these dialogs are the following container-level rights available:

- Create and place in folder
- Share/move from its parent folder
- Modify item behavior or configuration

The following table describes the access rights that are available from the Change Request, Requirement, Task, or Topic nodes in the *Project Access Rights*, *View Access Rights*, *Folder Access Rights*, and *<item> Access Rights* dialogs.

Table 12.8 <item> Access Rights

Generic Item Rights	
This access right...	**Allows a user or group to...**
See item and its properties	See the items in the upper pane and view item properties by selecting <item> > Properties.
Modify properties	Change item properties. The properties that can be modified are different for each type of item. Select <item> > Properties to review these properties, including custom properties. An alternate property editor (APE), if used, may restrict the properties that can be modified and the users who can modify them still further.
Delete from folder	Delete the items from their folders.
Change item access rights	Change the access rights for the items. If you change this setting, be sure that you remain one of the users who can change access rights.
See history	See the history for the items in the history pane.
Perform maintenance	Change the revision comment for past revisions.
Set exclusive locks	Lock the item exclusively.
Break exclusive locks	Remove someone else's exclusive lock on the items.
Label Rights	
This access right...	**Allows a user or group to...**
Attach/Adjust view labels	Add a view label to the items. Move a view label from one revision to another. This right controls direct manipulation of view labels for the items at the item level. It does not stop users from attaching a view label to the items when a view label is created.
Detach view labels	Remove a view label from the items. Be aware that if users can delete view labels, they can detach a view label from the items by deleting the view label from the view, regardless of the setting of this right.
Attach/Adjust revision labels	Add a revision label to the items. Move a revision label from one revision of the items to another. This right controls direct manipulation of revision labels for the items at the item level.
Detach revision labels	Remove a revision label from the items. Be aware that if users can delete revision labels, they can detach a revision label from the items by deleting the revision label from the view, regardless of the setting of this right.
Links Rights	
This access right	**Allows a user or group to...**
See links	See the links created for the items.
Create links	Link the items to other items.
Modify links	Change the properties of a link to the items. For example, you can change whether the ends are pinned or floating, and you can change the description.
Delete links	Delete a link that affects the items.

Table 12.8 <item> Access Rights (continued)

Generic Item Container Rights	
This access right...	**Allows a user or group to...**
Create and place in folder	Add items to a folder, view, or project if its *<Item> Access Rights* dialog has this option. This right appears only on the *<Item> Access Rights* dialogs associated with a folder, view, or project.
Share/Move out of folder	Share items or move items that reside in a folder, view, or project if its *<Item> Request Access Rights* dialog has this option. This right appears only on the *<Item> Access Rights* dialogs associated with a folder, view, or project. Be aware that the access rights set for any item that is moved or shared go with that item into the new folder.
Change behavior or configuration	Change the branching ability and configuration of items that reside in a folder, view, or project if its *<Item> Access Rights* dialog has this option. This right appears only on the *<Item> Access Rights* dialogs associated with a folder, view, or project.

Setting Access Rights Exceptions

Suppose that you have a group called Testers that has complete access to the files in the QA view, a view that contains folders full of test files. A newly hired member of the Testers group, New Tester, has not yet been trained to update the tests, etc. Although New Tester is a member of the Testers group, you do not want this user to perform certain operations on these files for a couple of weeks. You could remove New Tester from the Testers group temporarily, but the application also allows you to give New Tester all the rights of the Testers group with a few exceptions. To list the exceptions, you create a deny record.

To create an exception:

1. From an *Access Rights* dialog, click Add. The *Assign Access Rights To* dialog appears.
2. Select the user who is an exception within the group.
3. Select Deny, then click OK to return to the *Access Rights* dialog.

Caution Never select Deny to create an exception to a group unless that group is already specifically granted access for this same node. In this example, the Testers group must have access for this node.

4. Select and/or clear the appropriate check boxes.
5. Drag the deny record to the top of the list in the *Access Rights* dialog. Use the Move Up and Move Down buttons if you prefer.
6. Click OK.

Important All deny records must precede all grant records. Otherwise, the exception will not occur. For example, if the application finds the grant record for Testers before it finds the deny record for New Tester, the rights the user has as a member of the Testers group will apply.

Additional issues to be aware of include the following:

- Depending on the privileges of the Testers group, New Tester may be able to perform these operations anyway.
- If a deny record is the only record for a node, anyone not specifically granted access rights for that node has no access to that type of object at that level (unless ownership or privileges intervene). When the application finds a node for the correct

type of object that has even one record, it does not check higher levels for access rights.

Understanding Access Rights for Promotion States

Each view has its own set of promotion states. Access to these states is controlled by:

- The “Define promotion model” right, which is available from the View node of the *Access Rights* dialog at the view and project levels. See “Granting Access Rights at the View Level” on page 197.

  A user with the Define promotion level right can do anything to the promotion model:

  - Create and delete states.
  - Edit their properties.
  - Promote a label from one state to another. Promotion is a subset of editing properties. Anyone who can edit the properties of a state can also promote that state.
  - Reorder the states within the view.

- Access rights that govern access to individual promotion states. These Generic object rights and Promotion state specific rights are available from the Promotion State node of the *Access Rights* dialog at the view and project levels. They also appear on the access rights for individual promotion states.

  The rights for an individual promotion state are checked at the state level; if necessary, the checking continues at the view level and eventually the project level. If a user is granted a given right at one level, there is no need to check the next.

  When a right is granted at the view level, it applies to all states in the view, unless access is denied at the state level.

  When a right is granted at the project level, it applies to all the states in all the views within the project, unless access is denied at the state or view levels.

Table 12.9 Promotion State Access Rights

This access right...	Allows a User or Group to...
Change object access rights	Change the access rights for an individual promotion state. If you change this setting, be sure that you remain one of the users who can change access rights. This right is a generic object right. After creating a promotion state, you must exit and reenter the *Promotion* dialog if you want to set access rights for the newly created state.
Modify label assignment	Change the label assigned to an individual state either by clicking the Promote button or editing the label property. No other properties for the state can be edited unless the user also has the Define promotion model access right from the View node. This right is a promotion state specific right.

Setting Promotion State Access Rights

Setting promotion state access rights is very similar to setting other access rights. The access rights can be set at the project or view level as well as for individual promotion states.

To set promotion state access rights at the project or view level:

1 Select Project > Access Rights or View > Access Rights from the menu bar. The *Access Rights* dialog appears.

2 Select the View node.

3 From the *Access Rights* dialog, click Add to select a user or group. The *Assign Access Rights To* dialog appears.

4 Select a user or group. Users are listed by their user names and groups are listed by their paths (excluding the All Users group).

5 Select the Grant option button.

6 Click OK to return to the *Access Rights* dialog.

7 Select and/or clear the appropriate check boxes.

8 Click OK.

To set access rights for individual promotion states:

1 Select View > Promotion from the menu bar. The *Promotion* dialog appears.

2 Select a promotion state from the list box.

Note You can create the promotion state in the *Promotion* dialog. However, you must click Apply before you can set access rights. After you click Apply (or close and reopen the dialog), the Rights button is enabled.

3 Click Rights. The *Access Rights* dialog appears.

4 From the *Access Rights* dialog, click Add to select a user or group. The *Assign Access Rights To* dialog appears.

5 Select a user or group. Users are listed by their user names and groups are listed by their paths (excluding the All Users group).

6 Select the Grant option button.

7 Click OK to return to the *Access Rights* dialog.

8 Select and/or clear the appropriate check boxes.

9 Click OK.

Understanding Access Rights for Components, Filters, and Queries

Each filter or query resides in a particular component (such as the Change Request component or the File component) and can be applied to that component's type of data only in any project view managed by a specific server configuration.

Any user can create and use private filters and queries, but public filters and queries have access rights, individually and per component. Rights set on a specific filter or query take precedence over access rights set at the component level.

To apply a public filter or query, a user must be able to access the data type for the component in some open project view. When you apply the filter or query, it affects the type of data that visible in the open project view.

Users can apply any public filters and queries that they can view. In general, users can see any public filters and queries for which they are the owner (creator) or for which they have access rights. (Remember that the application can be set to ignore object ownership.)

Setting Component-level Access Rights

If you have the server-level access right to "Administer component-level access rights," you can set component-level access rights from any open component.

To set component-level access rights:

1 Open any project view to which you have access.

2 Select the correct tab for the component.

3 Select <component> > Advanced > Component Access Rights from the menu. The <*component*> *Access Rights* dialog appears.

4 Select the appropriate nodes:

- To control who can create public filters and queries for the component, use the Component node.
- To control who can use public filters for the component, use the Filter node.
- To control who can use public queries for the component, use the Query node.

5 Add a user or group.

a Click Add to display the *Assign Access Rights To* dialog.

b Select a user or group.

c Select the Grant option button.

d Click OK to return to the *Access Rights* dialog.

Caution Never select Deny unless you are creating an exception. See "Denying Access Rights" on page 200.

6 Select and/or clear the appropriate check boxes.

7 Click OK.

Note Clicking Delete removes the selected user or group from the User and Groups list. The selected user or group loses any previously set access rights to this component's filters and queries.

Table 12.10 Component Access Rights (At Component Level)

This access right...	Allows a user or group to...
Create public filters	Create public filter for this component.
Create public queries	Create public queries for this component.

Table 12.11 Filter Access Rights (At Component Level)

This access right	Allows a user or group to...
See object and its properties	See public filters for this component in the filters list (on the toolbar) and view their properties in the *Filters* dialog.
Modify properties	Change public filter properties for this component. The properties that can be modified for a filter are its list of displayed fields, its sorting and grouping rules, the query associated with it, and its context (the items of the component to which it can be applied).
Delete object	Delete public filters for this component from its list of filters.
Change object access rights	Change access rights for public filters for this component.

Table 12.12 Query Access Rights (At Component Level)

This access right...	Allows a user or group to...
See object and its properties	See public queries in the *Queries* dialog and view their properties in the *Edit Query* dialog.
Modify properties	Change public queries properties for this component. The properties that can be modified are the query's name and its conditions.
Delete object	Delete public queries for this component from its list of queries.
Change object access rights	Change the access rights for public queries for this component.

Setting Individual Filter and Query Access Rights

Filters and queries appear in every project view within the same server configuration. For example, if Project A and Project B are in the same server configuration, you can use the same filters and queries in each project view. The access rights for each query and filter is are same in all the views for each server configuration.

If users can see a view, they can use any of the server configuration's filters and queries in that view. If a filter or query is public and no access rights have been set for it, any user can apply that filter or query. If access rights have been set, only those users specifically granted access to that filter or query (as a user or as a member of a group) can see it and use it.

Table 12.13 Individual Filter Access Rights

This access right...	Allows a user or group to...
See object and its properties	See the filter in the filters list (on the toolbar) and view its properties in the *Filters* dialog.
Modify properties	Change the properties for the filter. The properties that can be modified for the filter are its list of displayed fields, its sorting and grouping rules, the query associated with it, and its context (the items of the component to which it can be applied).
Delete object	Delete the filter from the list of filters
Change object access rights	Change the access rights for the filter.

To set access rights for a filter:

1. Do one of the following:
   - Right-click a column header on upper pane, then select Filters from the context menu.
   - Select Filters > Filters from the appropriate menu or context menu.

   The *Filters* dialog appears.
2. Select the filter.
3. Click Access Rights. The *<filter> Access Rights* dialog appears.
4. Select a user or group. Users are listed by their user names and groups are listed by their paths (excluding the All Users group).
5. Select the Grant option button;
6. Click OK to return to the *Access Rights* dialog.

Caution Never select Deny unless you are creating an exception. See "Denying Access Rights" on page 200.

7 Select and/or clear the appropriate check boxes.

8 Click OK.

Note Clicking Delete removes the selected user or group from the User and Groups list. The selected user or group loses any previously set access rights to this filter.

Table 12.14 Individual Query Access Rights

This access right...	Allows a user or group to...
See object and its properties	See this query in the *Queries* dialog and view its properties in the *Edit Query* dialog.
Modify properties	Change the properties for this query. The properties that can be modified are its name and conditions.
Delete object	Delete this query from the list of queries.
Change object access rights	Change the access rights for this query.

To set access rights for a query:

1 Do one of the following:

- Right-click a column header on upper pane, then select Queries from the context menu.
- Select Filters > Queries from the appropriate menu or context menu.

The *Queries* dialog appears.

2 Select the query.

3 Click Access Rights. The *<query> Access Rights* dialog appears.

4 Select a user or group. Users are listed by their user names, and groups are listed by their paths (excluding the All Users group).

5 Select the Grant option button.

6 Click OK to return to the *Access Rights* dialog.

Caution Never select Deny unless you are creating an exception. See "Denying Access Rights" on page 200.

7 Select and/or clear the appropriate check boxes.

8 Click OK.

Note Clicking Delete removes the selected user or group from the User and Groups list. The selected user or group loses any previously set access rights to this query.

Using the Security Guidelines

Until you become familiar with access rights, Borland recommends that you follow the guidelines suggested in this section.

From the Server

On the Server, the User Manager dialog allows you to create users and groups for each server configuration while that configuration is running.

Use the following guidelines:

- Do not change the privileges for the All Users, Administrators, System Managers, and Security Administrators groups.
- Do not create additional groups under the Administrators group.

- Create the groups that you need under All Users or under each other. For example, you may need to create the following groups: Developers, Testers, and Writers.
- Create users and assign them to groups. Make sure that at least two users are administrators, in case one administrator becomes locked out.

From the application

Use the following guidelines:

- Although you can deny rights as well as grant them, it is best only to grant them.
- If you do deny rights, observe both of the following rules:
  - Never allow any node on an *Access Rights* dialog to have only deny rights records.
  - Always make sure that the deny rights records for any node precede any records that grant rights for that node.
- When you set access rights for a node, remember that any user who does not have access rights for the node (individually or in a group) is denied all rights at this level for this node (unless that user has ownership or privileges that allow access).
- Set access rights at the project level first. Set them for every group (except the All Users group) for every node. The nodes are Project, View, Child Folders, File, Change Request, Requirement, Task, and Topic. Depending on which version of the application your company uses, you may not see all of these nodes.

  The most important nodes to set at this level are the Project and View nodes. The Project node is the only location in which you can set project access rights. The View node controls view-level access to all views. Newly created views start out with only the view access rights set here for all views. Initially, they have no view-level access rights.

  The Promotion State node is discussed elsewhere. See “Understanding Access Rights for Promotion States” on page 212.
- Set access rights at the view level next. Set rights for every user and/or group that needs access at this level for every node. (The nodes are View, Child Folders, File, Change Request, Requirement, Task, and Topic).

  The Promotion State node is discussed elsewhere. See “Understanding Access Rights for Promotion States” on page 212.
- Set up access rights at the folder level only if you really need to have access rights for the folders. Remember that these settings go with the folder when it is moved or shared and when it becomes part of new views (until the folder branches in the new view). Remember that folders branch only when their properties change, and that their properties tend to change infrequently.
- Avoid setting access rights on root folders because those rights can conflict with those set at the project or view levels.
- Avoid setting access rights on items. Remember that these settings go with the item when it is moved or shared and when it becomes part of new views (until the item branches in the new view).

Chapter

13

Using StarTeam with Visual SourceSafe

When you are running the Server on a Windows platform, you can interoperate with Microsoft Visual SourceSafe (VSS) or convert VSS files to native StarTeam files.

When the application is first introduced at a company that uses VSS, some team members usually continue to use the VSS user interface and integration tools, while others rely on the StarTeam user interface to access VSS files. Both interfaces can be used with the same files, because StarTeam projects can be “wrapped around” VSS projects. You perform the wrapping process by creating a project in StarTeam that uses a VSS project as its root folder.

The application interoperates with VSS in three ways:

- Initially, a project is created by wrapping a set of VSS archives.
- As StarTeam users check file revisions into StarTeam, StarTeam updates the VSS archives so that VSS users can use the new revisions.
- As VSS users check file revisions into VSS, the application’s foreign refresh feature makes the new revisions available to StarTeam users.

Important The person running the Server must have the operating system privilege necessary to access the COM dll (SSAPI.dll) for VSS because the application uses the VSS Com API to perform file check-in and other operations.

Getting Ready

To prepare for wrapping, do the following:

- Install VSS on the computer that runs the application Server. Because the application uses the VSS database file (by default, srcsafe.ini) to locate the VSS files, the path to this database file and all the path names in that file must be recognized by the Server. You may need to use UNC names for these paths and for the project paths in the ss.ini files.
- Place the VSS projects to be wrapped on the same computer. Doing this will improve performance by reducing network traffic.

- Add the users in the VSS database to the server configuration that will manage the new project. The users' logon IDs in VSS and StarTeam must be identical (except for case). For any user not listed in the server configuration, the application will substitute the name of the user creating the project.

Note If VSS and the Server run on Windows systems and VSS users use their Windows logon names as VSS logon names, you can import these names into the Server. Doing this can save you time and prevent any typographical errors made when user names are added to the Server.

- Shorten VSS passwords if they are more than seven characters in length. Users who access the VSS files from the application must have VSS passwords that are no longer than seven characters. The VSS COM interface used by the application to access VSS will not accept passwords longer than this, although the VSS interfaces will.

  The Windows client asks each user for a VSS password. That VSS password is used until the Server is restarted or until VSS refuses it, in which case, the client asks the user for another password.

  If the user has either a blank VSS password or the same password in both VSS and the application, the Cross-Platform client can access the VSS archives.

- (Optional) Make sure all team members who will be using the application have checked in the tip revisions of their files, as the resulting folder structure will be strictly hierarchical. The working folder structure for these files may change significantly after the wrap.

- Undelete any files you wish to include in the wrap. VSS allows you to delete and undelete files, but the application does not wrap deleted files. Changes must not be made to the VSS files during the wrapping process.

- Check pinned files. A pinned revision is a revision that has been specified as the only revision of a file that can be checked out. If a pinned file has been shared in VSS, it should be pinned to the same revision in all locations. The application generates an error message during the wrapping process if one of these files is not pinned or if it is pinned to a different revision. As a result, it may or may not set the item configuration properties for these items correctly. These items are identified in an error message in the vss-err.*locale*.log file:

  ```
  File <VSSPathFile1> was shared to <VSSPathFile2> but they have different
  pinned times. The item configurations in StarTeam may not be set correctly.
  ```

  If a file included in the wrapping process is pinned after the wrap is finished or if a new file is pinned in VSS and recognized by the application via a foreign refresh, the file is treated as though the most recent revision is the pinned revision. This occurs because there is no way to tell when a file was pinned in VSS.

If your team decides to use StarTeam exclusively, you can convert the VSS data to the application data format after the wrapping process has taken place. See "Using StarTeam Exclusively" on page 226.

Understanding Branch Support

VSS files can contain named branches that represent parallel lines of development. During project creation, the application identifies a branched file as two separate files, one in the original VSS project or subproject and the other in the new VSS subproject, which coincidentally have identical history up until the time at which the branching occurred. No historical information is lost, but the fact that these two files were once the same file is not known to the application.

Wrapping the Application Around VSS

StarTeam enables you to create projects using VSS projects. It reads the VSS folder hierarchy and the entire revision history and lock status of each VSS file in the hierarchy. In addition, the application calculates the size of each revision.

Important You cannot wrap the same VSS data as part of more than one project. If you attempt this, the application may display the error message “Can’t get file creation date,” while VSS displays “A history operation is already in progress.”

The wrapping process may take a long time. It is greatly affected by the number of revisions and the folder structure of the VSS projects. If the VSS projects are on the same computer as the Server, network traffic is reduced, which saves time. However, if VSS is on the same computer as the Server, checking out files from VSS places the files in StarTeam cache. Doing this should be avoided because it makes the data about these files unreliable.

Caution VSS projects can be archived. If you archive data and delete it from the Repository, you may not be able to work with the project. If you encounter problems, verify the VSS version number on the computer where the Server is installed. If you have a version of VSS that is later than version 6, you need to locate a copy of VSS version 6 or earlier. Rename your existing SSAPI.DLL to SSAPIversionX.DLL. This action prevents you from overwriting the current DLL file. Copy the SSAPI.DLL file from VSS version 6 or earlier to your existing VSS installation location and register it on Windows using regsrv32. Then restart the Server.

Dates

During the wrapping process, if a VSS file has revisions with check-in dates that do not chronologically increase, the application modifies the check-in dates to make them chronological. The application logs the changes in the vss-err.*locale*.log file in the repository folder and displays a message telling you to examine the file before the wrapping process completes.

When creating derived views, the creation date given to each view is the time stamp of the oldest revision in the view.

Labels

VSS has two types of labels: project labels and revision labels. During the wrapping process, all project labels become view labels if they appear in any archives being wrapped. All revision labels in VSS become revision labels in the application. The application retains only the VSS labels it had at the time of the wrapping. If you convert to StarTeam immediately after wrapping, no label problems occur. If you interoperate for some length of time, you may prefer to re-wrap the project before converting its data to StarTeam formats.

In VSS, a project label is created by selecting a folder, clicking File > Label, and entering a label name and a comment. VSS applies the label to all the folders and files in the folder hierarchy that are rooted at this project. However, a child folder can apply the same label, and “override” the date and comment fields of the parent folder. Each uniquely named project label becomes a view label in the application. The time and the comment field for the highest label in the project hierarchy become the time and description for the view label. If a file has more than one project label with the same name (because the same label was applied to different folders), the application applies it to the most recent revision with that label.

Revision labels can be created during the check-in process or by using File > Label. Each label has a name and comment. These labels can be added, modified, or deleted from any revision at any time. VSS revision labels can be attached to only one revision of a file. If you revise a file and try to attach a revision label that has already been used with an earlier revision, VSS forces you to either cancel the operation or remove the

label from the earlier revision. A VSS revision label can have the same name as a VSS project label, but no application revision label can have the same name as an application view label. Therefore, the application will not create a revision label if a project label with the same name exists. The comment for a VSS revision label becomes the description for a StarTeam revision label.

In VSS, you can create a new version of a file with no changes in file contents by selecting the file and using File > Label to add a label. In this situation, the application generates a new version of the file to match the version in VSS. If desired, you can delete one of the file versions in VSS by adding the label a second time, which makes a gap in the VSS history. During the wrap the application recognizes the gap but cannot obtain any data, so it generates a revision one second after the most recent content revision. The application uses the revision comment in the file history to indicate that the revision is a "fake."

Pinned files

If a file is pinned in VSS prior to the wrap, the StarTeam item configuration for the file will show that it is pinned. Because VSS does not supply the time at which the pin operation occurred, the application uses a time immediately after the pinned revision was checked in. If a user rolls back a view to a configuration time that is after the time at which the application pinned the revision, the item will remain pinned.

Pinned files in both StarTeam and VSS are read only and therefore cannot be modified.

To wrap an application project around a VSS project so that the projects are interoperable:

1. Do one of the following:
   - Click the New Project icon.
   - Select Project > New from the menu bar.
   - Drag the srcsafe.ini file from Network Neighborhood or Explorer and drop it on an empty portion of the application window to create the project.
2. Do one of the following:
   - From the server list, select the server configuration that will store the project.
   - Click Add Server to add access to another server configuration and follow the steps in "Adding Access to a Server Configuration" on page 74. Then select the server configuration.
3. If you are not currently logged onto the server configuration, you will be asked to do so. After you log on, the *New Project Wizard* dialog appears.

   If you are already logged onto the server or are running the application Toolbar, but wish to log on as a different user, you can do so. On the *New Project Wizard* dialog, click the Log On As button and enter an alternate user name and password in the *Log On* dialog. To reset your default credentials for the server, select the "Save as default credentials for this server" check box . When you are done, click OK.
4. From the *New Project Wizard* dialog, click Next. The *New Project Wizard: Project Name* dialog appears.
   a. Enter the name of the new project in the Project Name text box.
   b. Enter a description of the project.
   c. If VSS security is disabled, go to step 5. If it is not disabled, go to step 4d.
   d. Click Foreign Archives.The *Select Foreign Archive* dialog appears.
      1. Select the Create Project From check box.
      2. Select the Visual SourceSafe option button.

3. In the Database File Path text box, enter or browse for the database file path (by default, srcsafe.ini).
4. In the Password text box, enter the VSS password for the user currently logged on to the application.
5. Click OK. This action re-displays the *New Project Wizard: Project Name* dialog.

5. From the *New Project Wizard: Project Name* dialog, click Next. This action displays the *New Project Wizard: VSS Project* dialog, which displays the VSS project hierarchy within the VSS database.
6. Select the VSS project or subproject from the list that will be the root folder for the new project, then click Next. The *New Project Wizard: Working Folder* dialog appears.
7. In the Default Working Folder text box, enter or browse for the path to the working folder for the root folder.

   The application wraps itself around the specified VSS project and its subprojects. It recursively creates working folders for each VSS project beneath the project selected as the root. To create working paths for each of these folders, the application starts with the path to the working folder for the root folder and appends folders for each level in the hierarchy, using the subproject names.

   However, the application does not support user-specified working folder paths except at the view level, which may make some or all of the VSS working folders obsolete.

   If any problems are found, an error message indicates that you should check the vss-err.*locale*.log file. If the vss-err.*locale*.log file problems are severe, you may want to delete the project and re-wrap the VSS projects.
8. Click Finish to open the project. Any VSS files with working folders that happen to match the working folders in StarTeam display in the file list with the Unknown status. You must update the status of these files on your workstation.

Security Issues

When you open a StarTeam project that wraps a VSS project, you are asked to log onto VSS when you perform an operation that accesses a VSS file.

If you access a VSS file from the Cross-Platform client, you will see a message indicating that a logon failure has occurred unless the VSS password is either blank or the same as your StarTeam password.

If you access a VSS file from the View Comparison utility, you will see a logon failure message. To rectify this problem, use the Windows client, open the wrapped project view, press F6, and log on. Then you can return to the View Comparison utility.

When the application prompts a user for a VSS password, that password is sent over the connection with the same level of encryption as other data. Therefore, it is important that the connection be encrypted on the server.

Note Sometimes, when the application requests that you enter a VSS password, it will ask you to repeat your last command.

VSS to StarTeam Interoperability

The following changes to the VSS database are recognized and/or reproduced in the application.

- Checking out a file (setting a lock).

- Undoing a file checkout (removing a lock).
- Adding revisions.
- Adding and deleting files.
- Adding subprojects (adding a folder).
- Modifying the working file following an add or check-in operation (when keyword expansion is enabled).

These changes in VSS are relayed to StarTeam.

- Sharing files and folders (VSS 6.0 or higher only).

When the following functions are performed in VSS, they are not currently recognized by or reproduced in StarTeam. Some of these operations can cause StarTeam to generate errors when accessing VSS.

- Deleting folders (however, you can delete the folder from StarTeam in a separate operation).
- Recovering deleted folders or files.
- Renaming projects or subfolders. The application treats the renamed project or folder as a new folder, which results in two copies of the folder in StarTeam, one with the old name and one with the new name. As a work around, you can delete the folder with the old name from the folder hierarchy. Never rename folders in StarTeam because VSS will continue to use the previous name.
- Moving projects. The application classifies a moved project or folder as a new folder, which results in two copies of the folder in StarTeam. Attempts to access the original folder in StarTeam produces errors in VSS because the folder is no longer in that location. As a work around, you can delete the folder in the old location from the folder hierarchy.
- Adding labels to revisions.
- Branching. The application does not read the action records in a file to re-create the branching hierarchy. Therefore, it classifies branched files as totally separate files that happen to have the same history up until the point at which they branched. No historical information is lost, but the fact that these two files were once the same file is not known to StarTeam.
- Historical changes to the folder hierarchy in VSS are not visible in rollback views in StarTeam.
- If you roll back a file in VSS, StarTeam does not recognize new revisions of that file until the number of VSS revisions for the file reaches its pre-rollback number.
- Changes to working folder paths in VSS are not recognized by StarTeam. In addition, working folder paths are not stored per user. Currently, the application forces working folder paths to be relative to the root folder for the view.
- Renamed files are recognized by StarTeam as new files. Operations on the original files will result in errors in StarTeam.
- Rolling back revisions is not currently implemented in StarTeam. If performed in VSS, you must delete the files that have been rolled back from StarTeam, then press F6 to re-synchronize with VSS. The deleted files will reappear, but with fewer revisions in their history. If you roll back the application view, you will see two files with the same name, but they may have different labels.
- Revisions that were pinned after the wrapping process took place are not recognized by StarTeam. A pinned revision has been specified as the only revision of a file that can be checked out. If revision 10 of 12 revisions is the pinned revision, the application considers this file to have the status Out-of-date, even though you cannot check out any revision other than 10.

- File descriptions that are updated in VSS are not updated in StarTeam. After a wrap, the description for each file is Initial revision. If you modify the description in VSS, and do a refresh in StarTeam, this description does not change.

None of the database options available in VSS are used in StarTeam. For example, the Delete working file on check-in option in VSS has no effect on StarTeam clients that access the VSS database. However, each StarTeam user can set some equivalent Personal Options in the application.

Of course, options set in the VSS Administrator project do affect the way in which data is added and retrieved from the VSS database. These options can be considered similar to the Server options. For example, the filter for keyword expansion is used to determine which files are processed for keyword expansion.

StarTeam to VSS Interoperability

The following functions in StarTeam result in equivalent operations in VSS:

- Checking in a file.
- Obtaining any revision of a file.
- Checking out a file.
- Locking a file (performs a check-out operation in VSS, but ignores the resulting file).
- Unlocking a file (performs an uncheckout operation in VSS).
- Breaking a lock when VSS security is disabled (performs an uncheckout operation for the user who had the file locked, followed by a check-out operation for the user who breaks the lock).

Note Security is considered disabled if the Administrator has the empty string as a password.

- Adding a file. If a deleted file with the same name exists in VSS, that file is undeleted.
- Adding a folder.

Because StarTeam is a true client-server application, all repository time stamps are translated to the time zone of the server, then forwarded to VSS. In this way, clients in different time zones can work on the same set of files and still maintain a valid chronology for the revisions. For this reason, it is important to keep the times on all workstations and servers accurate.

The following functions in StarTeam are not currently translated in VSS.

- Renaming a file. To rename StarTeam files, you must change the file name in the *File Properties* dialog. If the file is a foreign file, the Name text box is disabled.
- Renaming a folder.
- Sharing a file.
- Moving a file. Files from different VSS projects can be moved to the same StarTeam folder, but this operation should be avoided, as it is not translated into VSS.
- Branching a file.
- Representation of views. Neither reference nor branching views can be used in a StarTeam project created by wrapping a VSS project. If you want to use more than one view, you need to convert to StarTeam from VSS.
- Changing the working path for a folder does not change the working path in VSS, even if the StarTeam user is also logged into VSS.
- Linking. Although you can link items, the links are available only in StarTeam.

- Breaking a lock if VSS security is enabled that is, if the administrator has a password that is non-empty. To break a lock, you must log onto another client as the person who locked the item and then unlock the file.
- Changing the compression options of files in Native-I or Native-II format. These options do not appear for foreign files in the *File Properties* dialog. Deleting folders or files; however, a folder can be deleted from VSS as well.
- Forced check-ins of files that have not changed. These are usually rejected by VSS; however, in VSS 6.0, a forced check-in may be accepted about once per file.

Refreshing Data

StarTeam queries the VSS database about the current status of files before displaying them in a project view window. When you press *F6* (foreign refresh) or perform an operation on a StarTeam item, the application calls a synchronous routine. It uses a refresh feature rather than the event hook mechanism in VSS because the latter can cause the repository to become out of sync with VSS.

Note The foreign refresh command is available only from the Windows client.

The synchronous routine updates the current folder in StarTeam. Each existing file is checked for new revisions and lock status. In StarTeam, each new file revision must have a check-in date later than that of the previous revision. Because this restriction does not exist in VSS, StarTeam changes check-in dates on revisions when necessary.

StarTeam recognizes new VSS files and folders and deletes files that were deleted from VSS. However, it does not delete folders that were deleted from VSS.

By using rollback views, StarTeam can display the historical data for a deleted item. However, because the deleted item was stored in VSS rather than a StarTeam repository, StarTeam cannot recover revisions of the deleted item. Pressing *F6* does not refresh the screen in a rollback view, because these views are treated as read-only.

Before pressing *F6*, enable the All Descendants button so that the refresh recurses through all the known child folders.

Sometimes pressing the *F6* command displays an error message indicating that StarTeam is unable to refresh a specific item and that the user should delete it from StarTeam and then press *F6* again. When you follow these instructions, the current view will have a correct representation of the foreign archives. However, in a rollback view, the deleted item will appear twice. Both items reference the same foreign archive. If the first item is checked out, the status of the second item may become Merge or Unknown. It is not necessary to do a forced checkout of the second item, however, because the data has already been checked out to the same location.

Important If several users simultaneously do foreign refreshes, performance may become slower.

Using StarTeam Exclusively

Wrapping a StarTeam project around a VSS project makes it possible to use either product to access the files in the VSS project. After the wrapping process has been completed, you can permanently convert the VSS project to StarTeam. After the conversion, changes to the files will be reflected in StarTeam only. StarTeam will no longer wrap the VSS project, and you can delete the VSS database files.

To ensure that no new VSS subprojects, branches, or shared files are missed in the conversion, press *F6* until all folders and files available in VSS are also available in StarTeam. In the course of the conversion, folders and files in the folder hierarchy for the VSS project will be converted to StarTeam storage types.

Depending on the size of the VSS project, the conversion may be time-consuming. It takes as long as the initial wrapping process. Before beginning the conversion, you should make sure that all users are logged off both the StarTeam and VSS projects.

As noted earlier, archived VSS projects may cause problems during the conversion process. These problems can be resolved by using the SSAPI.DLL from VSS version 6 or earlier.

To convert an interoperable StarTeam project to StarTeam storage:

1. In StarTeam, open the project that wraps the VSS project.
2. Select Tools > Convert to Native Format from the menu bar. A message asks you to confirm the conversion of the files in this project to the StarTeam storage type.
3. Click OK. Another message indicates that the conversion process has begun.

   If e-mail notification is enabled, the application will send you (the logged-on user) an e-mail when the process completes. For more information, see “Enabling E-mail Support” on page 28. Depending on the number and size of the archives being converted, the process may take some time. .

A GoNative log file provides data about the conversion. It is stored in *repositoryFolder*\Logs\Files_Logs, where *repositoryFolder* is the location set for the server configuration repository. The log file has the name GoNative.*locale*.log.

To differentiate GoNative log files, time stamp information is added to the existing GoNative log file when you create a new one. The earlier file’s name changes from GoNative.*locale*.log to GoNative.*locale*.time.log. For example, you might see both GoNative.en_US.Log and GoNative.en_US.2004-02-18-14-38.Log in the Files_Logs folder.

Chapter

14

Using StarTeam with PVCS

If you run the Server on a Windows platform, you can interoperate with PVCS and/or convert PVCS files to native StarTeam files.

When StarTeam is first introduced at a company that uses PVCS, some team members usually continue to use the PVCS user interface, while others rely on the StarTeam user interface to access PVCS archives, the files that store the evolutionary history of each working file. Both interfaces can be used with the same files, because StarTeam projects can be “wrapped around” PVCS archives. You perform the wrapping process by creating a project in StarTeam for the PVCS archives.

The application interoperates with PVCS in three ways:

- Initially, a StarTeam project is created by wrapping a set of PVCS archives.
- As StarTeam users check file revisions into the application, it updates the PVCS archives so that PVCS users can use the new revisions.
- As PVCS users check file revisions into PVCS, the application’s foreign refresh feature makes the new revisions available to StarTeam users.

The application uses the PVCS command line to perform operations. In some cases, the command line operates differently from the PVCS graphical user interface.

Most companies prefer to convert PVCS files to StarTeam native files and stop using PVCS altogether. If you do not convert to native files immediately after wrapping the files, you should rewrap before the conversion. By doing this, you ensure that labels created after the original wrapping will be stored in the server configuration database.

Understanding the Terminology Differences

PVCS has project databases, each of which contains a hierarchical set of projects. A PVCS database references a configuration file (.cfg). In StarTeam, a PVCS.cfg file can be turned into a StarTeam project, which contains a hierarchical set of folders. In other words, the PVCS project database becomes a StarTeam project, and the PVCS projects become StarTeam folders. Subsets of a PVCS database can also be turned into StarTeam projects.

Assume that you have a PVCS project database (projectdb1) and three projects (bigproject, d1, and d2). When this database is mapped to StarTeam, it will contain a project named bigproject and three folders (bigproject, d1, and d2).

Each PVCS file is actually an archive from which past revisions of the file can be derived as working files. Although StarTeam stores files differently, it can also retrieve past revisions as working files.

PVCS archives may contain branches with floating labels that represent parallel lines of development, while StarTeam uses views for this purpose.

Understanding PVCS Configuration Files

To create a configuration (.cfg) file that works well with StarTeam, you may want to start with the .cfg file created by PVCS and add to it with Notepad or another text editor.

You should usually name the .cfg file in such a way that it identifies your project, such as “My Project.cfg.” If the file is not local to the server configuration, the path to it must be either on a permanently mounted disk or visible via a UNC file name (for example, \\machineX\pvcs project\My Project.cfg).

Look through your PVCS installation files to find the file called master.cfg, as it may contain information that you can paste into My Project.cfg.

After creating My Project.cfg, verify that it contains the following information.

- The most important directive in your PVCS file starts on the first line. It should be similar to the following:

  ```
  VCSDIR = "X:\My Project\xxx"; "X:\My Project\yyy"
  ```

  In this example, X: would be the drive letter of your PVCS archive, My Project would be the name of your PVCS project, and xxx and yyy would be folders that you intend to include in the wrap because they contain PVCS archive files.

  VCSDIR should include an entry for each folder that contains the PVCS archive files that are to be included in the wrap. Separate the entries with semi-colons (;).

  “My Project.cfg” must list the paths to the folders in which PVCS stores the project archives. Use path names that are relative to StarTeam Server. If a path in “My Project.cfg” lists “D:\....\....”, then the computer on which the Server runs must be able to access that location as “D:\...\...”. You can use UNC names, but they must still be relative to the Server.

- If you have access rights turned on in PVCS, a line might appear in master.cfg that looks like the following:

  ```
  ACCESSDB= "D:\Accessdb\Accessdb.db"
  ```

  You can copy this line and paste it into “My Project.cfg.”

- Add the following directives in the .cfg file:

  NOCASE=VCSID, because StarTeam logon IDs are not case sensitive.(The PVCS GUI does not add this directive to the vcs.cfg file.)

  NOMULTILOCK USER, because StarTeam allows only one user to lock a revision at a time.

- Remove the following directive from the .cfg file:

  EXCLUSIVELOCK, because this overrides NOMULTILOCK USER.

- PVCS refers to the .cfg file to determine the suffixes for archive files. The ArchiveSuffix directives provide the conversion information for the suffix of the PVCS files. Master.cfg may have the suffix lines already in place and you can just paste them into My project.cfg. It is important to paste all of them in each .cfg file. For example, your .cfg might contain directives similar to the following:

  ArchiveSuffix +-arcArchiveSuffix . +.-arc

  ArchiveSuffix = mak=.mvk

ArchiveSuffix =.bas=.bvs

ArchiveSuffix .asm=.avm

ArchiveSuffix = .bat=.bvt

These directives provide an extension and describe how it will look in the PVCS archive. PVCS replaces any single letter of the suffix with the letter “v.” The “v” can be anywhere, as in autoexec.vat, autoexec.bvt, autoexec.bav.

If you are wrapping solely to go native and not to interoperate with PVCS, you can use only the following ArchiveSuffix directive:

ArchiveSuffix = .???___

The long underscore above is really three underscore characters. The initial period is optional.

- StarTeam ignores any setting for the following directives:

  AUTOCREATE: the application always creates an archive before checking revisions in.

  CHECKLOCK: the application locks PVCS archives before checking files in, regardless of this directive.

  DELETEMESSAGEFILE: the application does not supply descriptions by using the contents of a file.

  DELETEWORK: the application deletes working files only if requested via StarTeam.

  FORCEUNLOCK: the application always allows unchanged revisions to be checked in.

Sample PVCS .cfg File

```
VCSDIR = "D:\WINNT\PVCSPRIV\sample\shifter\archives"
NOCASE VCSID
NOMULTILOCK USER
ACCESSDB= "D:\Accessdb\Accessdb.db"
ArchiveSuffix = .mak=.mvk
ArchiveSuffix = .bas=.bvs
ArchiveSuffix = .asm=.avm
ArchiveSuffix = .bat=.bvt
ArchiveSuffix = .doc=.dvc
ArchiveSuffix = .?=??v___
ArchiveSuffix = .ash=.avh
ArchiveSuffix = .cpp=.cvp
ArchiveSuffix = .cpr=.cvr
ArchiveSuffix = .dot=.dvt
ArchiveSuffix = .frm=.fvm
ArchiveSuffix = .frx=.fvx
ArchiveSuffix = .inc=.ivc
ArchiveSuffix = .ini=.ivi
ArchiveSuffix = .map=.mvp
ArchiveSuffix = .pbd=.pvd
ArchiveSuffix = .pbl=.pvl
ArchiveSuffix = .pbr=.pvr
ArchiveSuffix = .wav=.wvv
ArchiveSuffix = .xls=.xvx
ArchiveSuffix = .xlt=.xvt
ArchiveSuffix = .xlw=.xvw
```

Getting Ready

The wrapping process must be performed on a computer that has StarTeam installed and a hierarchy of archived PVCS files on its hard drive. Before you perform the wrapping process, you must prepare for this operation.

To prepare for wrapping:

- Install PVCS on the computer that runs the Server. The Server must be running, but not as an NT service.
- Verify that the user running the server process has execute privileges for the PVCS library, usually C:\Program Files\PVCS\VM\Win32\Bin. This can be checked by executing the vlog program. Depending on how PVCS was installed, this may involve modifying the user ID membership in particular user groups or the protection attributes on the PVCS installation files. See the PVCS Version Manager Installation Guide.
- Verify that the PVCS Bin folder is in the computer's path statement. Otherwise, neither the PVCS command line nor the StarTeam interoperation with PVCS will work properly. The PVCS installation does not set the PATH environment variable, so you must do so manually. Be sure that the environment variable is either set for all users or set specifically for the user who starts the server.
- Verify that the PVCS command line is working.
- Place the PVCS archives to be wrapped on the same computer as the Server. Although this is not required, it reduces network traffic and, therefore, improves performance.
- Rename any files with problematic paths prior to wrapping. StarTeam cannot process PVCS archives that contain parentheses in their file names or that have file paths with more than 128 characters. Some special characters, such as the backslash, are counted as two characters.
- Because StarTeam uses the PVCS configuration file (by default, pvcs.cfg) to locate the PVCS files, verify that the path to this file and all the path names within it are recognized by the Server. You can use UNC names, but they must still be relative to the Server.
- Verify that the directives listed in the configuration file are appropriate.
- Add the users in the PVCS archives to the server configuration that will manage the new StarTeam project. The users' logon IDs must be identical, except for case.

  For any revision author who is not listed as a user in the server configuration, StarTeam substitutes the name of the user creating the project. If an archive is locked by someone who is not found as a user in the server configuration and NOMULTILOCK is set, the application shows the file as unlocked but does not allow a StarTeam user to lock it.

  If PVCS and the Server run on Windows systems and PVCS users use their Windows logon names as PVCS logon names, you can import the user names into the Server, which saves time and prevents typographical errors.
- Check archives for labels. All labels must be unique. It is recommended that you use no more than one floating label per branch. All label names should be unique.

If your team decides to use StarTeam exclusively, you can convert the PVCS files to StarTeam files, stored in a StarTeam repository, after the wrapping process has taken place. See "Using StarTeam Exclusively" on page 239

Note Do not make any changes to the PVCS archives during the wrapping process.The wrapping process must be performed on a computer that has StarTeam installed on its hard drive.

Understanding Branch Support

PVCS archives can contain branches with floating labels that represent parallel lines of development. StarTeam uses views for this purpose. During the wrapping process, views can be created for branches that have floating labels. If a branch does not have a floating label, none of its descendants can become views in StarTeam, even if they have floating labels themselves. The wrapping process creates a view for the first, and only the first, floating label associated with each branch. The view name is the same as that of the floating label.

The StarTeam view is a branching view that is set to branch all files and has a floating configuration. (See Table 9.1, “Creating Views with Desired Characteristics,” on page 141 for more details.)

The additional labels result in error messages in the pvcs-err.*locale*.log file and should be removed from the file. For example, suppose an archive has two floating labels for the same branch. StarTeam will use the first label as a view name, ignore the second, and place an error message in the log file about the second. Suppose that the second label is also used by another archive as its first floating label on a branch. The result will be a StarTeam view with the second floating label’s name.

This situation will cause confusion when you interoperate with the first archive, because you are working with a file that does not exist in the view but has a label with the same name as the view. In this case, you will see an error message.

If you remove the label from the archive, the error message goes away and the problem is resolved.

Important The PVCS Developer’s Toolkit API allows a maximum of 64 characters for dot notation. When the dot notation for a revision exceeds 64 characters, it is truncated by the API. When this occurs, StarTeam skips the revisions having partial dot notations, because there is no way to correctly place the revisions within the archive. During the wrapping process, it alerts the user about this problem and logs the archive and revision numbers that were skipped.

A branch that does not have a floating label is reported in the pvcs-err.*locale*.log file. A message notifies you about this error. If the branch labels are important, you can stop the wrapping process, add floating branch labels where appropriate, and start the wrapping process again. If the branch labels are unimportant, you can continue the wrapping process, knowing that these branches will not become views in the StarTeam project.

If a floating label is on a trunk rather than a branch, it becomes a revision label on the last revision for the segment of the trunk to which it applies. For example, if a floating label moved from 1.1 to 1.2 to 1.3 when the numbering changed to 2.0, revision 1.3 receives the revision label.

If a PVCS archive does not contain the original branch label, StarTeam attempts to determine the branch to include in the view. It examines the archives that do contain the branch level to find the immediate ancestor of the branch. If the unlabeled archive contains the ancestor branch, StarTeam includes that branch in the view.

When different archives have different ancestors for the same branch label, StarTeam uses the most recent ancestor branch found in the unlabeled archive. That is the branch for which the associated PVCS revision number has the most decimals, which should be the most recent of the ancestor branches.

Fixed labels are generally converted to revision labels. The only exception occurs when an archive has a fixed label that matches the name of a view (that is, the fixed label has the same name as a floating label found in another archive). In this situation, StarTeam places the file in the view and pins its configuration to the revision with the fixed label.

Wrapping StarTeam Around PVCS

With StarTeam, you can create StarTeam projects using PVCS archives. The application reads the entire revision history and lock status for each PVCS file that it adds to the project. The application creates the project, then adds the views, folders, and revisions. After the project is created but before the process is finished, the project becomes accessible but should not be accessed by users.

During the wrapping process, if a PVCS archive has revisions with check-in dates that do not chronologically increase, the application modifies the check-in dates to make them chronological. StarTeam logs the changes in the pvcs-err.*locale*.log file in the repository folder.

Note StarTeam revision labels are different from PVCS fixed/revision labels. In StarTeam, a revision label belongs only to the view in which it was created. If a file has a given revision label in one view, it does not retain that revision label when it branches in a new view. However, when StarTeam wraps PVCS, a file that branches in a new view retains the revision labels it had in the parent view. This behavior is what PVCS users expect.

To wrap a StarTeam project around PVCS so that they are interoperable:

1. Do one of the following:
   - Click the New Project icon.
   - Select Project > New from the menu bar.
   - Drag the .cfg file from Network Neighborhood or Explorer to an empty portion of the StarTeam window to create the project.
2. Do one of the following:
   - From the server configuration list, select the server configuration that will store the project.
   - Click Add Server to add access to another server configuration and follow the steps in “Adding Access to a Server Configuration” on page 74. Then select the server configuration.
3. If you are not currently logged onto this server configuration, you will be asked to do so. After you log on, the New Project Wizard dialog appears.

   If you are already logged onto the server or are running the Toolbar, but wish to log on as a different user, you can do so. On the New Project Wizard dialog, click the Log On As button and enter an alternate user name and password in the *Log On* dialog. To reset your default credentials for the server, select the “Save as default credentials for this server” check box . When you are done, click OK.
4. Click Next. The *New Project Wizard: Project Name* dialog appears.
   a. Enter the name of the new project in the Project Name text box.
   b. Enter a description of the new project .
   c. If VSS security is disabled, skip to step 5. If it is not disabled, complete step 4d.
   d. Click Foreign Archives. The *Select Foreign Archive* dialog appears.
      1. Select the Create Project From check box.
      2. Select the PVCS option button.
      3. In the Config File Path text box, enter or browse for the .cfg file path.
      4. If security is enabled, enter the PVCS password for the user currently logged on to StarTeam in the Password text box.
      5. Click OK. This action redisplays the *New Project Wizard: Project Name* dialog.

5 Click Next. This action displays the *New Project Wizard: Working Folder* dialog.

6 In the Default Working Folder text box, enter or browse for the path to the working folder you want to use as the root working folder for your project.

7 In the Archive Root Folder text box, enter or browse for the root path to the archive folders. In your PVCS .cfg file, this is usually the first path name in the VCSDIR directive.

For example, on a Windows system, suppose that you want the working folder for your project to be C:\work and that you have the following VCSDIR directive in your .cfg file:

```
VCSDIR = "c:\archives"; "c:\archives\pc"; "c:\archives\mac"; "c:\archives\unix"
```

All of these paths have “c:\archives” in common and at the beginning of the path. Therefore, that would be the only correct choice for the archive root folder.

Also, if your paths contain spaces, they must be in quotation marks in the VCSDIR directive. Some directives use wildcards.

Suppose that you have the following VCSDIR directive in your .cfg file:

```
VCSDIR = "c:\archives"; "archives\pc"; "archives\mac"; "archives\unix"
```

No valid archive root folder is possible, so you must change this directive.

“c:\archives” would not work because it is not at the beginning of each path. You could use it, if you add c:\ to the paths that do not have c:\. You could use “\archives” if you removed the c: from the path that starts with those characters.

When StarTeam creates working folders for each archive folder, it replaces the archive root folder in each path in the VCSDIR directive with the path to the root working folder.

If you have the following VCSDIR directive in your .cfg file:

```
VCSDIR "c:\archives";"c:\archives\pc";"c:\archives\mac";"c:\archives\unix"
```

the result is the following working folders: C:\work, C:\work\pc, C:\work\mac, and C:\work\unix.

Note For the archive root folder on a Solaris system, you must use forward slashes (/) and remember that names are case sensitive. Working folders also require forward slashes.

8 Click Finish to open the project.

After the wrapping process is finished, the PVCS files display in the file list with Unknown status. You must update the status of these files on your workstation.

If any problems occur with the wrapping process, an error message asks you to check the pvcs-err.*locale*.log file. If the problems in the pvcs-err.*locale*.log file are severe, you can delete the project and re-wrap the PVCS archives.

Security Issues

When you use the Windows client to open a project that wraps PVCS archives, you are asked to log onto PVCS when you perform an operation that accesses an archive.

If you access a PVCS archive from the Cross-Platform client, you will see a message indicating that a logon failure has occurred unless the PVCS password is either blank or the same as your StarTeam password.

If you access a PVCS archive from the View Comparison utility, you will see a logon failure message. To rectify this problem, ruse the Windows client, open the wrapped project view, press F6 and log on. Then you can return to the View Comparison utility.

When StarTeam prompts a user for a PVCS password, that password is sent over the connection with the same level of encryption as other data. Therefore, it is important that the connection be encrypted on the server.

Note When StarTeam requests that you enter a PVCS password, it will sometimes ask you to repeat your last command.

PVCS to StarTeam Interoperability

The following changes to the PVCS database are recognized or reproduced in StarTeam.

- Adding a revision to an archive already known to StarTeam.
- Checking in a file.
- Locking and unlocking a tip revision.
- Adding new archives. When you press *F6*, new PVCS archives are added as files to a StarTeam view. In StarTeam child views, files are added only if a branch exists in the archive that has a floating label with the same name as the child view.
- Renaming an archive. The result in StarTeam is two files, one with the prior archive name and one with the new name. You can delete the file with the original name. Rollbacks are not possible after the archive has been renamed. If you try to check out a file from StarTeam that has been renamed in PVCS, you will see an error message.

When performed in PVCS, the following functions are not currently represented or reproduced in StarTeam. Performing some of these operations can cause the application to generate errors when accessing PVCS. For others, there are workarounds.

- Deleting a file. Files that are deleted from PVCS are not deleted in StarTeam. Otherwise, rollbacks are not possible. If you use the application to check out a file that has been deleted from PVCS you will get an error message.
- Adding a project to a PVCS project database. This operation is equivalent to adding a StarTeam folder to a StarTeam project; however, it must be done manually. When you add a project to PVCS, you add a folder for it to the StarTeam hierarchy.
- Changing the configuration file property for a PVCS project. This operation is not recommended, because the files in the PVCS project are considered by the application to belong to another PVCS database (that is, another StarTeam project).
- Changing the archive location for a PVCS project. This operation is not recommended, unless there are no files in the PVCS project. If you decide to set the new archive location property for a PVCS project, add that path manually to the VCSDIR directive in the .cfg file. Otherwise, this change will not occur in StarTeam.
- Promoting archives. StarTeam does not recognize any PVCS promotional model.
- Adding branches with fixed labels or no labels. (Fixed labels are ignored by StarTeam, except during the wrapping process when they become revision labels.) However, if a PVCS user applies a fixed label to all archives and branches several of them, a new view can be created in StarTeam with the same name as the fixed label. Pressing *F6* will populate the new view with the branched revisions of archives with that label.

  The view must be a branching view that is set to branch on change and that has a floating configuration.
- Adding branches with a floating label. StarTeam recognizes new PVCS branches only after you create a StarTeam view with the same name as the floating label for the PVCS branch. This view must be a branching view that has a floating configuration and is set to branch on change.

- Changing a file description in PVCS. These changes do not go to StarTeam. After a wrap, the description for each file is “Initial revision,” and it remains that regardless of changes in PVCS.

StarTeam to PVCS Interoperability

The following functions in StarTeam result in equivalent operations in PVCS:

- Adding a file. Files cannot be added directly to a child view in StarTeam. Instead, they must be added to the root view, and then automatically shared into the child view and set to branch on change. If you want a file to exist only in the root view, you must explicitly delete it from each child view before it branches in that view. When a file branches in a child view, it branches in PVCS and acquires a floating label for that branch.

  Never delete a file from the root view or from a child view that has its own child views. If a file is deleted from a view and not deleted from its child views, *F6* restores the file to the view and creates a second copy of the file in each child view.

  Files that you add via StarTeam are immediately accessible as PVCS archives from the PVCS command line. However, from the PVCS graphical user interface, you *must* use the Import Archive command to make the archive visible.
- Checking in a file.
- Getting any revision of a file.
- Checking out a file.
- Locking and unlocking a file in the root view or in a child view in which the file has branched. If a file has not branched in the child view and you lock it from the child view, PVCS is unaware of the lock.
- Unlocking the file for a former user and relocking it for the current user (that is, breaking a lock) if PVCS security was disabled when the server started. Security is considered disabled if the user who locked the file has the empty string as a password.
- Creating views/branches. Because views created in StarTeam become branches in PVCS, the application does not allow a PVCS file to be branched into more than one view with the same name. To work with PVCS branches, child views created in the application must have the same root folder as their parent views.
- Turning views into branches with floating labels. StarTeam does not create a branch in the PVCS archive for a file unless the file’s behavior is “Branch on Change” and a revision of that file has been checked into the new view. The appropriate revision is branched, but revision numbers in PVCS and in StarTeam may differ. Consequently, reference views never appear in the PVCS archives and should be avoided in projects that wrap PVCS archives.

Because StarTeam is a true client-server application, all repository time stamps are translated to the time zone of the Server and then forwarded to PVCS. In this way, clients in different time zones can work on the same set of files and still maintain a valid chronology for the revisions. It is very important to keep the times on all workstations and servers accurate.

The following functions in StarTeam are not currently translated in PVCS.

- Renaming a file. To rename StarTeam files, you change the file’s name in the *File Properties* dialog. If the file is a foreign file, the name text box is disabled.
- Deleting a file. If you delete a file from a view via StarTeam, PVCS does not know that it has been deleted. If you delete a file with StarTeam before a revision was checked in for a particular view, it does not have a branch in the PVCS archive for the view from which it was deleted. In this case, to restore the deleted file is not a

simple Add File operation. You must open both the view from which you deleted the file and its parent view, then share the file from the parent view to the child view. These actions restore the file and its history.

- Sharing a file.
- Changing the working path for a folder.
- Locking and unlocking a file in a child view in which the file has *not* branched. If you lock a file from a child view in which it has *not* branched, PVCS is unaware of the lock.
- Linking. You can link items, but the links are available only in StarTeam.
- Labeling views or files. StarTeam labels are not translated in any way into PVCS.
- Changing file descriptions. StarTeam can create a PVCS file description as it adds a file to the project, but it does not record any changes made to that description.
- Breaking a lock if PVCS security is enabled. Security is considered enabled if the user who locked the file had a non-empty password when the server configuration started. To break a lock, log into another client as the person who locked the item and then unlock the file.
- Changing the compression options of files in Native-I or Native-II format. These options do not appear for foreign files in the *File Properties* dialog.

Refreshing Data

When PVCS users access the archives, the information about those archives becomes out of date in StarTeam. For example, if a PVCS user adds a revision to an archive, that revision does not immediately appear in the application. The same is true for lock status. To update this data, StarTeam uses the refresh feature instead of the event trigger mechanism in PVCS, because a PVCS user can suppress the event trigger and cause the information to stay out-of-date.

Operational Refresh

When a StarTeam user performs an operation (such as a file check-in) on a file, the application updates a single out-of-date archive actually, the branch of the archive that corresponds to the current view. If the file is locked by a PVCS user who is not a StarTeam user, the application cannot update the archive and notifies the StarTeam user that the file is locked.

Folder Refresh

Pressing *F6* updates the current StarTeam folder and its files, regardless of their PVCS location. It is recommended that you enable the All Descendants button before pressing *F6* so that the refresh recurses through the child folders. Files in rollback or read-only views are not refreshed.

When you refresh a folder, each existing file is checked for new revisions and lock status. StarTeam requires each new revision of a file to have a check-in date later than the previous revision. Because this restriction does not exist in PVCS, StarTeam changes the check-in dates in the StarTeam history for revisions when necessary.

New PVCS archives are added as files to a StarTeam view if they have a PVCS branch name that corresponds to a StarTeam view name. However, files deleted from PVCS are not deleted in StarTeam.

If a file is locked by a PVCS user who is not also a StarTeam user, the application does not show the lock, and pressing *F6* does not update the archive.

Important

- The foreign refresh command (*F6*) is available from the Windows client only. Sometimes pressing this command displays an error message indicating that StarTeam is unable to refresh a specific item and that the user should delete it from StarTeam and then press *F6* again. When you follow these instructions, the current view will have a correct representation of the foreign archives. However, in a rollback view, the deleted item will appear twice. Both items reference the same foreign archive. If the first item is checked out, the status of the second item may become Merge or Unknown. It is not necessary to do a forced checkout of the second item, however, because the data has already been checked out to the same location.
- When you delete a PVCS file, be sure to delete it from all locations in StarTeam and from the PVCS archive. If you delete a PVCS file from a StarTeam view that has child views and do not delete the file from the PVCS archive folder, using foreign refresh (*F6*) in the parent view on a Windows client will reinstate the file. If the file has branched, using *F6* in a child view produces two copies of the file and causes the following error message to appear:

  ```
  StarTeam has detected that this folder has at least 2 items matching the foreign archive...
  ```

Using StarTeam Exclusively

Wrapping a StarTeam project around PVCS archives makes it possible to use either StarTeam or VSS to access the files in the VSS project. After the wrapping process has been completed, you can permanently convert the VSS project to StarTeam. After the conversion, changes to the files will be reflected in StarTeam only. StarTeam will no longer wrap the VSS project, and you can delete the VSS database files.

All working files in the current StarTeam folder hierarchy on the computer on which you perform the conversion will be converted to StarTeam storage types. Make sure you have a working file in the hierarchy for each PVCS archive to be converted.

Depending on the size of the VSS project, the conversion may be time-consuming. It takes as long as the initial wrapping process. Before beginning the conversion, you should make sure that all users are logged off both the StarTeam and VSS projects.

To convert a PVCS interoperable project to StarTeam storage:

1. In StarTeam, open the wrapped project.
2. Select Tools > Convert to Native Format from the menu bar. A message asks you to verify the conversion of the files in this project to the StarTeam storage type.
3. Click OK. Another message states that conversion process has begun.

   If e-mail notification is enabled, the application will send you (the logged-on user) an e-mail when the process completes. For more information, see “Enabling E-mail Support” on page 28. Depending on the number and size of the archives being converted, the process may take some time.

A GoNative log file provides data about the conversion. It is stored in *repositoryFolder*\Logs\Files_Logs, where *repositoryFolder* is the location set for the server configuration repository. The log file has the name GoNative.*locale*.log.

To differentiate GoNative log files, time stamp information is added to the existing GoNative log file when you create a new one. The earlier file’s name changes from GoNative.*locale*.log to GoNative.*locale*.time.log. For example, you might see both GoNative.en_US.Log and GoNative.en_US.2004-02-18-14-38.Log in the Files_Logs folder.

Appendix

A

Reporting on Licenses

The application can create a license usage report that enables you to keep track of the number of users and products at their site.

You must have the appropriate server access rights to view the license usage information. See “Setting Server-level Access Rights” on page 190 for more information.

Running the License Report

The License Report provides licensing compliance information. This information may also be used by Borland Technical Support when responding to customer inquiries. The report can be run from either the application or the Server Administration tool.

To run the License Report:

1. Display the *Server Administration* dialog by doing one of the following:
   - From the computer on which the Server is installed, select Start > Programs > StarTeam > StarTeam Server *x.x* > Server Administration.
   - From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration (available with custom installations only).
   - From the computer on which the Server is installed, enter the following at the command line.

     x:\<*path to StarTeam Server*>serveradmin
2. Select a server from the server configuration list. If you have not yet logged in, you will be asked to do so. If you are using a client, you can administer remote servers only.
3. Click the License Usage icon in the toolbar. The system will display the following message:

   ```
   Pleaseforwardthesefiles"X:\...\UsageInfo.txt"and"X:\...\UsageInfo.sig"to
   Borland Software Corporation as requested
   ```
4. Click OK.

Understanding Licenses

The Server can be licensed for different sets of features. The features of your server are determined by its license. The features that a client can access on that server are determined by the client license, unless the client is licensed for more features than the server. In that case, the client can access no more than the server has to offer.

We recommend using the same level of licensing for both your clients and your server. For example, if you use workflow on a server, it is usually because you want everyone to use the special forms and workflows you have designed for one or more components. If you allow clients with other types of licenses to access this server, they will not be able to use the specialized workflow-based forms that have been created for a given component.

Appendix

B

Initialization Files

Initialization files have different locations on different Windows platforms. On NT, C:\winnt\Profiles is the *pathPrefix*. On Windows 2000, XP, and 2003, the *pathPrefix* is C:\Documents and Settings.

- ClientLicenses.st lists is located at \All Users\Application Data\Borland\StarTeam\ClientLicenses.st. This file is installed by StarTeam Runtime and the application clients.

  If the ClientLicenses.st file is missing, you are asked to register the product.

- ConnectionManager.ini (used for starting the application) is located at \All Users\Application Data\Borland\StarTeam\ConnectionManager.ini. This file is installed by StarTeam Runtime and the application clients.

  If the ConnectionManager.ini file is missing or corrupted, the application asks if you want it recreated. Reinstallation can also recreate the missing ConnectionManager.ini file.

- starteam-servers.xml lists the server configurations for which you have created (or will create) server descriptions. These descriptions are used while opening or creating projects. This file is located at *user*\Application Data\Borland\StarTeam\ServerList and is installed by StarTeam Runtime and the application clients.

- The Server's starteam-server-configs.xml file, which is used for server session information, is located at *Server InstallationFolder*\ starteam-server-configs.xml.

- starteam-client-options.xml is located at *pathPrefix**user*\ Application Data\Borland\StarTeam\starteam-client-options.xml. This file is installed by StarTeam Runtime and the application clients.

  The starteam-client-options.xml file lists personal option settings and any alternate working folders you have set with the application. You may want to back up this file or put it under version control.

  If the starteam-client-options.xml file is missing, the application automatically recreates it. However, the recreated file contains only the default settings for the personal options and no alternate working folder information.

  If the starteam-client-options.xml file is corrupted, you can delete it, but you may be able to edit it. See "Understanding ConnectionManager.ini" on page 244 for more information.You can tell that the starteam-client-options.xml file is missing or corrupted when:

- Your personal options are no longer correct.
- Changes you made to personal options disappear when you restart the application.
- Files do not change even though you have checked them out (because they have been copied to the wrong working folders).
- The application reports that old files are missing and does not see new files, because it looking for them in the wrong place.

Understanding ConnectionManager.ini

The ConnectionManager.ini file contains information that the client must be able to locate in order to run. It is created at the time that the application is installed. The following is an example of a ConnectionManager.ini file. In actual ConnectionManager.ini files, the x's are replaced by hexadecimal numbers.

```
[ConnectionManager]
WorkstationID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
```

Understanding starteam-server-configs.xml

The starteam-server-configs.xml file contains session options for one or more server configurations. Server session options specify the core information that the Server requires to start a server configuration. One starteam-server-configs.xml file exists per computer and is located in the same folder as the Server application. On Windows platforms, this file is usually located in C:\Program Files\Borland\StarTeam Server.

The session option information for each server configuration begins with the name of the configuration in brackets and is followed by a set of options and their settings. The Server creates and maintains this file, which is created when the first server configuration is created. The file is updated whenever a server configuration is created, modified, deleted, started, or stopped. Do not edit this file directly.

Borland recommends that you backup the starteam-server-configs.xml file or put it under version control.

Information in this file includes:

- CipherName. In server configurations created prior to StarTeam Server 5.3. For internal use. Do not edit this option.
- CipherSource. For internal use. Do not edit this option.
- CipherTest. For internal use. Do not edit this option.
- ChangeRequestsCaching. Caches tip revisions of change requests on the server. Defaults to 2 so that check-out operations are faster.

  Settings can be:

  0 Off

  1 Cache each tip revision when it is accessed by a user

  2 Cache all tips from all views
- Command. For internal use. Do not edit this option.
- CompressionOffFileExt. Applies to Native 1-files only. When this option is set to file specifications, the Server turns off compression (file storage) for files matching these specifications if they also exceed the size limitation set on CompressionOffFileSize. When blank, compression is turned off for all files.

You can use wild cards or enter multiple extensions separated by commas, spaces, or semicolons. To include a comma, space, or semicolon as part of the specification, enclose the specification in double quotation marks. For example, you might use:

```
CompressionOffFileExt=*.zip; *.avi "Big File.rtf"
```

By default, compression is turned on. You can adjust the compression setting inside the application on a per-file basis. However, for some files, you may prefer to have the Server identify files that match the limitations set with CompressionOffFileExt and CompressionOffFileSize and turn off compression automatically. See also CompressionOffFileSize, DeltaStorageOffFileExt, and DeltaStorageOffFileSize.

- CompressionOffFileSize. Applies to Native-I files only. When this option is set to a number of bytes (for example, 300 MB is 300000000), the Server automatically turns off file compression for files of this size or greater that match the file specifications set in CompressionOffFileExt. When this option is blank or set to zero, no files have compression turned off.

  See also CompressionOffFileExt, DeltaStorageOffFileExt , and DeltaStorageOffFileSize.

- ComputerName. ComputerName identifies the computer on which the Server resides. This option is set by the Server. Do not edit it.
- ConfigurationName. The server configuration's name appears in square brackets at the beginning of the server configuration entry. This information is set using the -new option with the starteamserver command. See “Using the starteamserver command” on page 253 for more information.
- CreatedByBuild. Indicates that a server configuration cannot have any Native-I archive files. Used only in server configurations created with StarTeam Server 2005 or later releases.
- DBCreated. DBCreated indicates whether the database tables used by the application are already created. Do not edit this option.
- DBServerName. Because all databases are accessed via ODBC, this is the DSN name for that database. In releases 5.1 and 5.2, Oracle databases were accessed using the Oracle net service name that is stored in $ORACLE_HOME/network/admin/tnsnames.ora. This is not the case for StarTeam Windows Server 5.3 or later.

  This information is set using the -dsn option with the starteamserver command. See “Using the starteamserver command” on page 253 for more information. You can review the DSN using the -`view` and -`edit` options from the command line or in the application on the Database tab of the dialog. Any modifications you make will take effect the next time you start the server configuration.

- DBType. Use one of the following numbered values to indicate the type of database.This information is set using the -t option with starteamserver command. See “Using the starteamserver command” on page 253 for more information. The database type can only be specified when creating a new server configuration. This information cannot be modified for existing server configurations.

  - 1 = Microsoft Access
  - 2 = MSDE or Microsoft SQL Server
  - 3 = Oracle
  - 6 = DB2 (not yet available)

  You can review the database type using the -`view` option from the command line or in the application on the Database tab of the dialog.

- DBUserName
  DBPassword

DBUserName and DBPassword are the name and password by which the application accesses the server configuration database.

This information is set using the -u (for the username) and -p (for the password) options with the starteamserver command. See "Using the starteamserver command" on page 253 for more information.

- DeltaStorageOffFileExt. Applies to Native-I files only. When this option is set to file specifications, the Server turns off delta storage for files that match these specifications if they also exceed the size limitation set using DeltaStorageOffFileSize. When blank, delta storage is turned off for any files.

  You can enter wild cards or multiple extensions separated by commas, spaces, or semicolons. To include a comma, space, or semicolon as part of the specification, enclose the specification in double quotation marks. For example, you might use:

  ```
  DeltaStorageOffFileExt=*.zip; *.avi "Big File.rtf"
  ```

  By default, delta storage is turned on in Native-I vaults. You can adjust the delta storage setting inside the application on a per-file basis, but, in the case of large files, you may prefer to have the Server watch for files that match the limitations set with DeltaStorageOffFileExt and DeltaStorageOffFileSize and turn off delta storage automatically.

- DeltaStorageOffFileSize. Applies to Native-I files only. When this option is set to a number of bytes (for example, 300 MB is 300000000), the Server automatically turns off file compression for files of this size or greater size that match the file specifications set in DeltaStorageOffFileExt. When blank or zero, no files have their delta storage turned off.

  See also DeltaStorageOffFileExt, CompressionOffFileExt, and CompressionOffFileSize.

- FilesCaching. Caches tip revisions of files on the server. Defaults to 2 so that check-out operations are faster.

  Settings can be:

  0 Off

  1 Cache each tip revision when it is accessed by a user

  2 Cache all tips from all views

- Initialized. Initialized indicates whether this server configuration was initialized. This option is maintained by the Server. Do not edit it.

- ListenIP. This option binds a server configuration to a specific TCP/IP (sockets) address. For example, if the Server has more than one IP address (more than one NIC card), you can configure the Server to listen to one specific port.

  When this option is set to 0 (default) the server configuration listens to all IP addresses on the specified port. (The port is specified on the Protocol tab of the dialog.)

- ListenXML. This option binds a server configuration to a specific IP (XML) address. For example, if the Server has more than one IP address (more than one NIC card), you can configure the Server to listen to one specific port.

  When this option is set to 0 (default) the server configuration listens to all IP addresses on the specified port. (The port is specified on the Protocol tab of the dialog.)

- MaxCommandThread. This option specifies the maximum number of command threads that the server configuration can create. The default maximum number of command threads is 16. (A setting of 0 for this option causes the server configuration to use the default value.) See the description of MinCommandThreads for more information.

MaxCommandThreads can be modified by editing starteam-server-configs.xml. However, we recommend that you change this option only when directed to do so by a Borland representative.

- MinCommandThreads. This option specifies the minimum number of command threads that the server configuration will create. When the server configuration starts and while it is running, it will have at least this many command threads. The default minimum number of threads is 4. (A setting of 0 for this option causes the server configuration to use the default value.) See the description of MaxCommandThreads for more information.

  MinCommandThreads can be modified by editing starteam-server-configs.xml. However, we recommend that you change this option only when directed to do so by a Borland representative.

- PID. PID is the process ID for the instance of the server configuration that is currently running. Otherwise this option is set to 0. This option is maintained by the Server. Do not edit it. When this option is missing, starteamserver creates it.
- RepositoryPath. RepositoryPath is the complete path to repository folders. This information is set using the -r option with starteamserver command. The repository path can only be specified when creating a new server configuration. This information cannot be modified for existing server configurations.
- Response. For internal use.
- ServerGuid. ServerGuid is a value supplied by the Server. Do not edit this option.
- UserName. UserName is the (domain) user name for the user who created the server configuration. This option is set by the Server. Do not edit it.
- ServiceMode. ServiceMode is for use on Windows NT systems only. Use 1 to run the server configuration as an NT Service. Use 0 to run the server configuration as an application.
- Status. Status indicates whether this server configuration is Ready, Starting, Running, or Stopping. This option is maintained by the server. Do not edit it. When this option is missing, starteamserver creates it.
- TasksCaching. Caches tip revisions of tasks on the server. Defaults to 0.

  Settings can be:

  0 Off

  1 Cache each tip revision when it is accessed by a user

  2 Cache all tips from all views

- TopicsCaching. Caches tip revisions of topics on the server. Defaults to 2 so that check-out operations are faster.

  Settings can be:

  0 Off

  1 Cache each tip revision when it is accessed by a user

  2 Cache all tips from all views

- VaultConversion. Controls the background process that converts Native-I archive files to Native-II archive files.

  Settings can be:

  0 Process is not running. (Default)

  1 Process is running (or will be when the server configuration starts).

- VaultConversion Mode

  Indicates the status of the archive files.

  Settings can be:

  -1 Conversion completed; all archive files now are Native-II archive files.

  1 Conversion not yet completed; when Native-I archive files are converted they are not deleted. By default, converted archive files are placed in a ConvertedArchives folder (repository_path\Archives\ConveredArchives).

  2 Conversion not yet completed; when Native-I archive files are converted, they are deleted.

Understanding starteam-client-options.xml

The starteam-client-options.xml file contains one line for each of the options that can be set from the *Personal Options* dialog (accessed by selecting Tools > Personal Options from the menu). Most option names in the starteam-client-options.xml file correspond closely to the names of the options on the dialog. The options that are check boxes in the dialog are set equal to 1 for selected or 0 for cleared. Intervals are set to a number of minutes or seconds depending on the option. Paths are in text. No quotation marks are used with the text.

For example, the Project Component information provides the paths to alternate working folders for projects accessed from your workstation. The entry for this component in the starteam-client-options.xml file includes the following parts.

- The phrase Project Component.
- ViewWorkingFolderOverrides (if the alternate working folder is for an entire view) or WorkingFolderOverrides (if the alternate working folder is for an individual folder).
- A hex identification of the project view and folder.
- The alternate working folder’s path.

Appendix

C

Backing Up the Application

To properly back up project data, it is necessary to back up all data at the same time. Projects consist of a database, archive files, and configuration files. If a mismatch between the database and the archive files occurs, the backup may be unusable. Borland recommends that a tape drive or similar backup device be used to back up application projects.

Important Be sure to verify your backups periodically. Borland recommends that you restore and test backups of your project data on a test system. Doing this helps to ensure that your data is being backed up correctly.

When you create server configurations, the Server creates files and folders that store your configuration and project information. For example, in StarDraw, the sample configuration that ships with the application, the Server places the hive, which contains Archives and Cache subfolders, and the Attachments folder in the Repository. The StarDraw Repository also contains a Database folder that holds MSDE database information (that is, the .mdf file). However, in other server configurations, the database files will be in a separate location.

Although your server configuration repository may differ, the underlying folder structure will be similar. You can change the location of its folders, if necessary, with the *StarTeam Server Configuration* dialog or the *Hive Manager* dialog for the Server.

What to Back up

The following files should be backed up.

- starteam-server-configs.xml

  This file contains the server configurations. It is located in the folder in which you installed the Server.

- Database files

  Each server configuration has one database (or, in the case of Oracle, one schema user).

- File revisions (repository)

  Hive\Archives folder: *.*

- Files attached to change requests, etc.

Attachments folder: *.*

All of these files should be backed up at the same time, preferably on the same tape.

How to Perform Back ups

Starting with StarTeam Server 2005, you have a choice between online and offline backups.

Online Backups

If all of your archive files are in Native-II format, you can back up a server configuration online—without shutting it down or locking it. If you have a pre-2005 StarTeam Server release, you must convert older archives from Native-I to Native-II format prior to performing an online backup. You must also set up your database for online backups. For releases prior to 2005, it was not necessary to run the database in archivelog mode (Oracle) or full backup mode (sql server). For more information about online backups, see he *StarTeam Installation Guide*.

The vault should **not** be backed up until after the database backup completes. The two backups should **not** be done simultaneously. In this way, you guarantee that everything referenced in the database appears in the vault in the right data location. The fact that the vault may contain files that the database does not know about causes no problems.

The HiveIndex.xml and the Attachments folder should be included in the vault backup, along with the archive folders from every hive. Backing up the cache folders for the hives is optional. Usually, the database and the vault are on different computers. The vault itself may be spread over several volumes and on different computers. For more information on this subject, see Chapter 7, Understanding Data Locations, in the *StarTeam Installation Guide*. This chapter explains what repository files needs to be backed up and what files can be ignored. For example, if you are using VSS or PVCS wrapping, you don't need the VSS_Hive or PVCS_Hive, but you must back up the actual VSS or PVCS archives.

In the event that the database is lost, the administrator must restore the last full backup of the database and apply the redo logs (Oracle) or apply incremental backups (sqlserver) to roll forward the database to the vault time.

In the event that the vault is lost, it is very important to take a backup of the database in its current state, including the transaction logs (redo logs), before doing any restoring.

In the event that the vault (or both the vault and database are lost), the admin must restore both the database and the vault from the last backup. After restoring the online database backup, the database has to be rolled forward to the vault backup time.

Note If you do not have transaction logs (redo logs) available, this can cause data loss and limit your disaster recovery capabilities.

Offline Backups

If you want to backup an archive that contains files in Native-I format and keep the files in that format, you can perform an offline backup.

Before running the backup, you must shut the server down by doing any of the following:

- If you are running the server configuration manually, the server configuration can be shut down in several different ways, including the following:

- From the client, select Start > Programs > StarTeam > StarTeam client_name > Server Administration. Select the appropriate server configuration, and click Shut Down. Only remote servers can be administered from the client.
- From the computer on which the Server is installed:
  - Select Start > Programs > StarTeam Server x.x > StarTeam Server and select the appropriate server configuration. Then click the Shut Down Local Server Icon or select Actions > Shut Down Server from the menu.
  - Enter the following at the command prompt:

    ```
    starteamserver -stop ConfigurationName
    ```

    where *ConfigurationName* is the name of the server configuration.
- If you are running the server configuration as a Windows service, go to the Services applet and stop the service. After the backup is completed, you may restart the service.

After the server is shut down, you can back up the database and repository.

Appendix

D

Using the starteamserver command

This appendix explains the options available with the starteamserver command and gives examples of their uses. Many of the functions performed by the starteamserver options can also be performed using the Server Administration utility. See Chapter 3, “Configuring and Running the Server,” on page 9 for more information.

The starteamserver command is available only on the computer on which the Server is installed.

StarTeamServer.exe (StarTeamServer.x on Solaris) is located in the folder that contains the Server and is run from the command line. The starteamserver command enables you to:

- Start, stop, and restart a Server.
- Create and modify the server configuration session options. The session options are stored in the starteam-server-configs.xml file. Most modifications do not take effect until the next time you start the server configuration.
- List information about the Server such as licensing and version information.
- Register the Server.
- View server configurations and their status.

The command syntax is as follows:

```
starteamseserver -option Value -option Value ...
```

The command options can appear in any order. If the option values contain spaces, be sure to enclose the value with double quotes; for example:

```
starteamserver -new NewServer1 -r "c:\new server\" -t 1 -dsn NewServerDSN -u
Admin -p password
```

The following table presents an overview of the tasks you can perform using the `starteamserver` options. For detailed information about the syntax and use of these options, see “Understanding starteamserver Options” on page 257.

Table 14.1 starteamserver Command Options

This option...	Enables you to...
`-start`	Start a server configuration with or without overrides. See “Starting a Server Configuration” on page 256.
`-stop`	Stop one or all running server configurations.

Table 14.1 starteamserver Command Options (continued)

This option...	Enables you to...
`-edit`	Edit the server configuration name, database user ID and password, or database connectivity information. These changes take effect the next time you start the server configuration.
`-restart`	Stop and restart the specified server configuration. Use this option after you make changes to a server configuration that you want to take effect immediately.
`-new`	Create a new server configuration. See "Creating a New Server Configuration" on page 254 for instructions.
`-serial access`	Register the Server as a fully licensed version.
`-eval`	Extend the evaluation period for the Server.
`-remove`	Delete a server configuration from the starteam-server-configs.xml file. You must manually delete the associated repository files.
`-help`	Display help for the starteamserver command.
`-version`	Display the build and version number of the Server.
`-licenses`	Display the license information for the Server.
`-view`	View the server configuration definition in the starteam-server-configs.xml file.
`-list`	List the status of all server configurations. A server configuration can have the status of Ready, Running, Starting, or Stopping

Creating a New Server Configuration

You can create server configurations using the command line or Server Administration. All server configurations created with StarTeam 2005 or a later release will use a Native-II vault.

To create a new server configuration from the command line:

1 Open a command prompt window and navigate to the server's installation folder.

2 At the command prompt enter the following command.

```
starteamserver -new "ConfigurationName"
-r "RepositoryPath" -t DBType -dsn "DataSourceName"
-u "DBUserName" -p "DBUserPassword"
```

where

- *ConfigurationName* is a unique server configuration name.
- *RepositoryPath* specifies the folder and files that the Server will create for this server configuration. The Server must be able to access this location.

Cautions

RepositoryPath must *not* be located in the server installation folder.

If you select a repository path that has been previously used by another server configuration, you will overwrite that server configuration's repository files. You must manually delete or move these files *before* using the new server configuration.

The starteamserver command puts the log file (Server.*locale*.Log) at this location. It also creates the following objects under the *RepositoryPath*:

- Server log files

  A new server log file is created every time you start the server configuration.

- Attachments folder

Attachments has child folders that store files attached to specific types of items. For example, the Change_Attachments folder stores files attached to change requests. Never change the names of this folder.

- HiveIndex folder

  This folder stores the hive-index.xml file, which contains the properties for each hive used by the server configuration.

- VSS_Hive or PVCS_Hive folder
- The application automatically creates VSS_Hive and/or PVCS_Hive folders in the repository path if it locates a VSS or PVCS installation where the Server is also installed.
- DefaultHive folder

  If you accepted all the defaults when you created the server configuration or if you started an upgraded server configuration without first creating a hive, the Server automatically creates the Default Hive folder as a subfolder in the *RepositoryPath*. Whether or not the initial hive is called DefaultHive, you will have at least one hive for each server configuration. The hive contains two subfolders, generally named Archives and Cache.

After the server configuration is in use, additional objects will probably be added under the *RepositoryPath.* For more details, see the chapter titled "Understanding Data Locations" in the *StarTeam Installation Manual*.

- *DBType* specifies the database type this server configuration uses. The database type can be specified only when creating a new server configuration.

  Use one of the following values to indicate the database type:

  - 2 = MSDE or Microsoft SQL Server
  - 3 = Oracle
  - 6 = DB2 (not yet available)

- *DataSourceName* specifies the database source name (DSN) created for the database. The name must appear between double quotes (") in the starteamserver command. The DSN must already exist.

  In releases 5.1 and 5.2, Oracle databases were accessed using the Oracle net service name that is stored in $ORACLE_HOME/network/admin/tnsnames.ora. This is no longer the case.

Caution *Never* create more than one server configuration that uses the same database. The table information for a server configuration will become corrupted if two different server configurations update the same database.

- *DBUserName* specifies the user name the Server uses to access the database.
- *DBUserPassword* specifies the password the Server uses to access the database.

3. The Server displays the following message when it finishes executing the starteamserver command.

   Configuration *ConfigurationName created successfully.*

   The system adds the new server configuration to the starteam-server-configs.xml file.

4. Start the server configuration by typing the following command:

   starteamserver -start "*ConfigurationName*"

The first time you start a new server configuration, the Server performs a number of startup tasks, including:

- Creating and initializing the server configuration's database.

- Installing the stored procedures for that particular database type.
- Creating the repository folders.

This process may take several minutes. When the Server is finished, it displays the following message.

```
Server ConfigurationName started successfully.
```

At this point you should review and update the server configuration options using the Server Administration utility. See "Setting Server Configuration Options" on page 21 for more information.

Starting a Server Configuration

You can start a server configuration using the values defined for it in the starteam-server-configs.xml and server configuration database or you can override these values, as explained below.

Note You can override certain server configuration values with the -restart option.

To start a server configuration using the defined values:

1. Open a command prompt window and navigate to the server's installation folder.
2. At the command prompt, enter the following:
   `starteamserver -start` `"`*`ConfigurationName`*`"`

To start a server configuration and override the defined values:

1. Open a command prompt window and navigate to the server's installation folder.
2. At the command prompt, enter the following:
   `starteamserver -start` `"`*`ConfigurationName`*`" [`*`options`*`]`

   where `[`*`options`*`]` can be any combination of the following. See "Understanding starteamserver Options" on page 257 for more details about these and other options.

   - `-name "`*`ConfigurationName`*`"`
     Renames the server configuration.
   - `-tcpip` *`Endpoint`*
     Specifies a new TCP/IP endpoint.
   - `-archive "`*`ArchivePath`*`"`
     (Applies only to Native-I vaults) Changes the location of the Archive folder.
   - `-attach "`*`AttachmentsPath`*`"`
     Changes the location of the Attachments folder.
   - `-cache "`*`CachePath`*`"`
     (Applies only to Native-I vaults) Changes the location of the Cache folder.
   - `-vault "`*`VaultPath`*`"`
     (Applies only to Native-I vaults) Changes the location of the Vault folder.
   - `-verifyvault`
     (Applies only to Native-I vaults) Verifies the tip revision of each file in the archive with the database each time you start the server configuration. Use this option with one or both of the following options:
   - `-autorecover`
     Instructs the Server to attempt to make limited repairs when necessary during the verification process.
   - `-stoponerrors`
     Instructs the Server to stop the startup process if errors are detected during verification. All errors are written to the server log file.

- `-noverifyvault`
  (Applies only to Native-I vaults) Turns vault verification off.
- `-dsn "`*DataSourceName*`"`
  Specifies the data source name (DSN) for the database.
- `-u "`*DBUserName*`"c`
  Changes the user name used by the server configuration to connect to the database.
- `-p "`*DBUserPassword*`"`
  Changes the password for the user name used by the server configuration to connect to the database.

Understanding starteamserver Options

This section describes the options for the starteamserver command in alphabetical order, with examples of their uses.

- `-access` *Key*

  Use with: `-serial`

  See also: `-serial`, `-license`, `-eval`

  Registers the Server as a licensed version. Use this option with the -serial option. The first time you start the Server, you must register the application as either a licensed version or an evaluation copy. If you need a serial number/access key combination or an evaluation key to extend your evaluation period, contact www.borland.com/company/contact/where_to_buy.html.

  Example:

  ```
  starteamserver -serial 1234 -access 5678
  ```

- `-all`

  Use with: `-start`, `-stop`, `-restart`

  Used in conjunction with the `-start` (or `-restart`) or `-stop` options. The `-start -all` options start all server configurations that have a status of Ready in the starteam-server-configs.xml file. The `-stop -all` options stop all server configurations that have a status of Running.

  Example:
  ```
  starteamserver -stop -all
  ```

`-archive "`*ArchivePath*`"` (Applies to Native-I vaults only)

Use with: `-start`, `-restart`

Specifies the archive path for a server configuration and is used in conjunction with the `-start` or `-restart` option.

With a Native-I vault, the system creates an Archive folder under the Vault folder and stores the path to this folder in the database used by the server configuration. If you change the location of the Archive folder, you can modify the archive path in the database using the `-start` (or `-restart`) and `-archive` options from the command line.

You can also modify the archive path on the Native-I vault tab of the *StarTeam Server Configuration* dialog. The new path takes effect the next time you start the server configuration.

Example:

```
starteamserver -start MyServer
-archive "c:  y Server\Vault\Archive"
```

- **-attach** "*AttachmentsPath*"

  Use with: -start, -restart

  Specifies the attachments path for a server configuration.

  The first time you start a server configuration, the system creates an Attachments child folder under the path you specify for *RepositoryPath* and stores the path to this folder in the database used by the server configuration. If you change the location of the Attachments folder, you can modify the attachments path in the database using the -start (or -restart) and -attach options from the command line.

  You can also modify the attachments path on the General tab of the *StarTeam Server Configuration* dialog. The new attachments path will take effect the next time you start the server configuration.

  Example:

  ```
  starteamserver -start MyServer
  -attach "c:\My Server\Attachments"
  ```

- **-autorecover**

  Use with: -start, -verifyvault

  See also: -stoponerrors

  Used in conjunction with the -verifyvault option, which applies only to Native-I vaults. The -autorecover option instructs the Server to attempt to make limited repairs where necessary during the verification process.

  You can also set this option in the application on the Native-I vault tab of the *StarTeam Server Configuration* dialog.

  Example:
  ```
  starteamserver -start MyServer -verifyvault -autorecover
  ```

- **-cache** "*CachePath*" (Applies to Native I Vaults only).

- Use with: -start, -restart

  Specifies the disk cache path for a server configuration and is used in conjunction with the -start or -restart option.

  For a Native I Vault, the system creates a Cache folder under the Vault folder and stores the path to this folder in the database used by the server configuration. If you change the location of the Cache folder, you can modify the archive path in the database using the -start (or -restart) and -archive options from the command line.

  You can also modify the Cache path on the Native-I vault tab of the *StarTeam Server Configuration* dialog. The new path takes effect the next time you start the server configuration.

  Example:

  ```
  starteamserver -start MyServer
  -archive "c:  y Server\Vault\Archive"
  ```

- **-dsn** "*DataSourceName*"

  Use with: -new, -edit, -start, -restart

  See also: -t, -p, -u

  Specifies the database connection information. Enter the existing ODBC data source name (DSN).

  In releases 5.1 and 5.2, Oracle databases were accessed using the Oracle net service name that is stored in $ORACLE_HOME/network/admin/tnsnames.ora. This is no longer the case.

The value you specify for *DBServerName* is stored in the starteam-server-configs.xml file. You can review or modify the database connection information by using:

- The `-view` and `-edit` options from the command line.
- Database tab of the *StarTeam Server Configuration* dialog in StarTeam Administration.
- Database tab of the *<server configuration> Properties* dialog in Server Administration.

Modifications take effect the next time you start the server configuration.

Example:
```
starteamserver -edit MyServer -dsn MyServerDSN
```

- **`-edit "`***ConfigurationName*`"`**

  Use with: -name, -dsn, -u, -p

  Edits the session options for the specified server configuration. You can edit the following options: `-name`, `-dsn`, `-u`, `-p`. If the server configuration is running, you must shut it down before you can make any edits.

  Example:

  ```
  starteamserver-editMyServer-namePortable-dsnRemoteServer-uStarTeamAdmin
  ```

- **`-eval`** *Number*

  See also: `-serial`, `-access`, `-license`

  Extends the evaluation period for an evaluation copy of the Server. The first time you start the Server, you must register the application as either a licensed version or an evaluation copy. If you need a serial number/access key combination or an evaluation key to extend your evaluation period, contact www.borland.com/company/contact/where_to_buy.html.

  Example:
  ```
  starteamserver -eval 01234567890
  ```

- **`-help`**

  Displays a message describing all of the command options.

  Example:

  ```
  starteamserver -help
  ```

- **`-licenses`**

  See also: `-serial`, `-access`, `-eval`

  Displays license and registration information. If you are running a evaluation copy of the application, the system displays a message informing you of this. Otherwise, the system displays your serial number.

  Example:

  ```
  starteamserver -licenses
  ```

- **`-list`**

  Lists the server configurations defined in the starteam-server-configs.xml file and the status of each one. A server configuration can have one of the following statuses at any given point in time; Ready, Starting, Running, Disabled, and Stopping.

  Example:
  ```
  starteamserver -list
  ```

  The Server displays a message similar to the following:

```
Configuration    Status
MyServer         Ready
StarDrawRepository Running
Portable         Ready
```

- **-name** "*ConfigurationName*"

  Use with: -edit, -start, -restart

  Renames a server configuration. This option is used in conjunction with the -edit option. The new server configuration name will take effect the next time you start the server configuration.

  Example:

  ```
  starteamserver -edit MyServer -name NewTeamServer
  ```

- **-new** "*ConfigurationName*"

  Creates a hive named DefaultHive for the new server configuration with the specified name and settings. This configuration uses a Native-II vault. This option produces the same result as selecting New on the *Server Administration* dialog and using the wizard to create a new configuration.

  A number of options can only be specified with -new, these are: -c, -r, and -t.

  Example:

  ```
  starteamserver -new NewServer1 -r "c:\new server\"
  -t 1 -dsn NewServerDSN -u Admin -p password
  ```

- **-noverifyvault** (Applies to Native-I vaults only)

  See also: -verifyvault

  .Disables vault verification.

  Example:

  ```
  starteamserver -start MyServer -noverifyvault
  ```

- **-p** "*DBUserPassword*"

  Use with: -new, -edit, -start, -restart

  See also: -dsn, -p, -t, -u

  Specifies the password used to access the database. The value you specify for *DBUserPassword* is stored in the starteam-server-configs.xml file. Ensure that the password you specify is the correct one for the database user name. You can review or modify the password and user name using the -view and -edit options from the command line. Any modifications you make will take effect the next time you start the server configuration.

  Example:

  ```
  starteamserver -edit MyServer -u JodyK -p password
  ```

- -q

  Executes a command in quiet mode—that is, the system displays no screen output.

  Example:
  ```
  starteamserver -stop -all -q
  ```

- **-r** "*RepositoryPath*"

  Use with: -new

  Specifies the repository path for a new server configuration. If the repository path you specify does not exist, the system will create the appropriate folders the first time you start this server configuration.

The value you specify for *RepositoryPath* is stored in the starteam-server-configs.xml file. You can review the repository path using the -view option from the command line or in the application on the General tab of the *StarTeam Server Configuration* dialog.

Important Do not use the Server home folder/directory as a server configuration repository path because the server configuration will not start.

Example:

```
starteamserver -new NewServer1 -r "c:\new server\"
-t 1 -dsn NewServerDSN -u Admin -p password
```

- **-remove** "*ConfigurationName*"

  Deletes the specified server configuration from the starteam-server-configs.xml file.

  Example:

  ```
  starteamserver -remove MyServer
  ```

- **-restart** "*ConfigurationName*"

  Stops and restarts the specified server configuration. Use this option after you make changes to a server configuration and want those changes to take effect. If the server configuration fails to restart, check the server log file for more information.

  You can restart a server configuration and modify a number of its options at the same time. The following options can be used with the -restart option: -all, -archive (Native-I vaults only), -attach, -cache (Native-I vaults only), -dsn, -name, -p, -tcpip, -u, -vault (Native-I vaults only), -verifyvault (Native-I vaults only). You cannot use both the -all and the specific configuration name at the same time.

  Example:

  ```
  starteamserver -restart MyServer -tcpip StarTeamTCPIP -u SuperUser -p
  SuperUserPassword
  ```

- **-serial Number**

  See also: -access, -license, -eval

  Registers the Server as a licensed version. Use this option with the -access option. The first time you start the Server, you must register the application as either a licensed version or an evaluation copy. If you need a serial number/access key combination or an evaluation key to extend your evaluation period, contact www.borland.com/company/contact/where_to_buy.html. The serial and access numbers in the example below would be replaced with actual serial and access numbers.

  Example:

  ```
  starteamserver -serial 1234567890 -access 9999999
  ```

- **-start** "*ConfigurationName*"

  See also: -all, -stop

  Starts the specified server configuration. starteamserver updates the server configuration entry in the starteam-server-configs.xml file to Status=Running and PID=*nnn* where *nnn* would be replaced with the actual PID number.

  You can start a server configuration and modify a number of its options at the same time. The following options can be used with the -start option: -archive (Native-I vaults only), -attach, -cache (Native-I vaults only), -dsn, -name, -p, -tcpip, -u, -vault(Native-I vaults only), -verifyvault (Native-I vaults only).

  Example:

  ```
  starteamserver -start MyServer -tcpip StarTeamTCPIP -u SuperUser -p
  SuperUserPassword
  ```

- -stop "*ConfigurationName*"

  See also: -all, -start

Shuts down the specified server configuration. After the server configuration stops running, starteamserver updates the entry in the starteam-server-configs.xml file to Status=Ready and PID=0.

Example:
```
starteamserver -stop MyServer
```

Note for Enterprise Advantage Users

If you are running the Server as a service and Notification Agent as a dependent service, you cannot shut down the Team Server unless the Notification Agent service is shut down first.

- **`-stoponerrors`** (Applies to Native-I vaults only)

  Use with: `-start`, `-verifyvault`

  Used in conjunction with the `-verifyvault` option (Native-I vaults only). Use `-stoponerrors` if you want the Server to stop the startup process if errors are detected during the verification process. All errors are written to the server log file.

  You can review and modify this option in the application on the Native-I vault tab of the *StarTeam Server Configuration* dialog.

  Example:
  ```
  starteamserver -start MyServer -verifyvault -stoponerrors
  ```

- **`-t`** *DBType*

  Use with: `-new`

  See also: `-dsn`, `-p`, `-u`

  Specifies the database type. This option can be used only when you are creating a new server configuration. Use one of the following numbered values to indicate the type of database:

  2 = MSDE or Microsoft SQL Server

  3 = Oracle

  6 = DB2 (not yet available)

  The value you specify for *DBType* is stored in the starteam-server-configs.xml file. You can review the database type using:

  - The `-view` option from the command line
  - In StarTeam Administration on the Database tab of the *StarTeam Server Configuration* dialog
  - In Server Administration on the Database tab of the *<server configuration> Properties* dialog

  Example:
  ```
  starteamserver -new NewServer1 -r "c:\new server\" -t 1 -dsn NewServerDSN -u Admin -p password
  ```

- **`-tcpip`** *Endpoint* `|` `up[:`*Endpoint*`]` **|** `down[:`*Endpoint*`]`

  Use with: `-start`, `-restart`

  Sets the endpoint for the TCP/IP (Sockets) protocol. Also enables or disables the protocol. Use `up` to enable and `down` to disable. You can both set the endpoint and enable or disable it using up or down followed by a colon and the endpoint.

  The value you specify for the endpoint is stored in the database used by this server configuration.

  You can modify this information using the `-start` (or -restart) and `-tcpip` options from the command line or in the application on the Protocol tab of the *StarTeam Server Configuration* dialog.

  Example:

```
starteamserver -start MyServer -tcpip 49201
starteamserver -start MyServer -tcpip up
```

- **-u** "*DBUserName*"
  Use with: -new, -edit, -start, -p, -restart
  See also: -t, -dsn, -p

  Specifies the user name that the server configuration uses to access the database. The value you specify for DBUserName is stored in the starteam-server-configs.xml file. You can review or modify the database user name using the -view or -edit options from the command line. Be sure to also specify the password for this user account. Any modifications you make will take effect the next time you start the server configuration. Ensure that the user name and password you specify using the starteamserver command is a valid account in the database. The server configuration will fail to start if the user account is missing in the database.

  Example:
  ```
  starteamserver -edit MyServer -u SuperUser
  -p SuperUserPassword
  ```

- **-vault** "*VaultPath*" (Applies to Native-I vaults only)
  Use with: -start, -restart
  See also: -verifyvault,-stoponerrors,-autorecover

  Specifies the vault path for a server configuration.

  If you are using a Native-I vault, when you first start a server configuration, the system creates a Vault child folder under the path specified for *RepositoryPath* and stores the path to this folder in the database for the server configuration. If you change the location of the Vault folder, you can modify the vault path in the database by using the -start (or -restart) and -vault options from the command line.

  You can also modify the vault path in the application on the Native-I vault tab of the *StarTeam Server Configuration* dialog. The new vault path will take effect the next time you start the server configuration.

  Example:
  ```
  starteamserver -start MyServer
  -vault "c:\My Server\Vault"
  ```

- **-verifyvault** (Applies only to Native-I vaults)

  Use with: -autorecover, -start, -stoponerrors

  See also: -noverifyvault

  Used in conjunction with the -autorecover and -stoponerrors options. This option instructs the Server to verify the tip revision of each file in the archive with the database each time you start the server configuration.

  Borland recommends verifying the database *only* after restoring a backup or when experiencing database problems. Most administrators *do not* routinely verify the database at startup.

  You can review and modify this option in the application on the Native-I vault tab of the *StarTeam Server Configuration* dialog.

  Example:
  ```
  starteamserver -start MyServer -verifyvault -stoponerrors -autorecover
  ```

- **-version**

  Displays the version and build number for the Server.

  Example:

  starteamserver -version

  The Server displays a message similar to the following:

```
StarTeam Server
Version: x.x
Build number: x.x.xxx
```

- **-view** "*ConfigurationName*"

  Lists the session properties of the specified server configuration.

  Example:
  ```
  starteamserver -view StarDraw
  ```

Appendix

E

Using the Server Administration Utility

The new Server Administration utility included with StarTeam 2005 Release 2enables you to work with the databases and vaults associated with server configurations. You can use it to:

- Upgrade a server configuration and the database it uses from an earlier version of the application. See the *StarTeam Installation Guide* for more information.
- Upgrade a Native-I vault to a Native-II vault (StarTeam 2005 or later releases).
- Launch background conversion of archive files in Native-I format to Native-II format (StarTeam 2005 or later releases).
- Create and manage hives in a Native-II vault (StarTeam 2005 or later releases).
- Verify that the files stored in a Native-I vault are recognized by the database as revisions of specific files. (StarTeam 6.0 or earlier).
- Purge deleted views from a server configuration database and vault.
- Migrate a server configuration from one database to another. Server Administration allows you to migrate from any database supported by the Server to any other. For example, you can use the Migrate feature to migrate an Oracle database to a Microsoft SQL Server 2000 database.
- Export catalog tables and fields to a comma-separated text file for use in a spreadsheet.

  Often the database table and field names are different from those used by the Server. The Catalog Export tool enables you to compare these names and other data and to translate database tables and column names into identifiers that are used by the Server.

Upgrading to StarTeam 2005 or Later Releases

Server configurations created with StarTeam 6.0 or earlier use Native-I vaults for storage, but configurations created with StarTeam Server 2005 or later releases use the new Native-II format. After you upgrade a server configuration created with StarTeam 6.0 or earlier to StarTeam 2005 or later releases, the configuration uses the Native-II vault for all new files put under application control. However, it continues to use the existing Native-I vault for existing files until they are converted manually or as part of a background conversion process.

New server configurations created with StarTeam 2005 and later releases use only Native-II vaults.

Native-II vaults store each revision of a file in a separate archive file, which allows a revision to be much larger than before. Compression is used for files that compress well (10% or better), but delta storage of archives has been eliminated. This format speeds up revision retrieval.

Native-I vaults store all revisions for the same file in the same archive file. While delta storage and compression can reduce the size of this file, it continues to grow with additional revisions.

Eventually, the Server will support only Native-II vaults.

The two types of vaults have many similarities. Native-I and Native-II vaults both contain an Attachments folder, the server log files, and other server-related items. Both vaults store .dmp files and trace files, if they exist.

However, the basic structure of the vaults differs.

A Native-I vault has only one Vault folder, a subfolder of the repository path. The Vault folder has one Archive and one Cache folder. Initially these folders are subfolders of the Vault folder but you can change their locations later, if necessary, with the Server Administration utility. For example, if the disk becomes full, you may need to move one or both of these folders (and their files) to create space for additional files and revisions.

A Native-II vault can have several hives, each of which has its own Archives folder and Cache folder. If one hive fills up, you can add another without having to change any data locations or move any archive files. When a server configuration has more than one hive, the server adds files to each of them in turn before reusing the first hive's archive path.

To upgrade a server configuration created with StarTeam 6.0 or earlier, see the *StarTeam Installation Guide*.

Creating New Hives

Native-II vaults store each file revision in its entirety (even though the archive file may be compressed). But the revisions can be spread over many volumes by the use of hives for storage. If one hive fills up, you can add another, without changing any data locations or moving any archive files. When a server configuration has multiple hives, the server adds files to each hive in turn before reusing the first hive's archive path. For more information about Native-II vaults, see the *StarTeam Installation Guide*.

When a server configuration is created with StarTeam 2005 or a later release, it automatically has at least one hive (either the default or a custom hive). To increase the amount of available space for this server configuration, one or more new hives can be added with the Hive Manager. Hives can be created while the server configuration is running, because the configuration already has an initial path, if only to a Default Hive in the repository path.

Note If you are adding a hive because the original hive was low on space, you should also use the *Hive Manager* to display properties the properties of that hive and deselect the "Allow new archives" check box. This action allows the original hive to remain a check-out location, but keeps it from acquiring any new files. Files that are added go to the new hive.

To create the initial hive for a server configuration created with StarTeam 6.0 or earlier, the application offers two different methods:

- Upgrade the server configuration to StarTeam 2005 or a later release, then use the *Hive Manager* to create a custom hive for that server, and then start the server. Once the server is started, it creates the Custom Hive, as well as the Native-II vault and its other files and folders. This method allows you to give the hive a unique name and to select your own location for the Native-II vault and the hive’s Archives and Cache folders; therefore, Borland recommends it.
- Upgrade the server configuration to StarTeam 2005 or a later release and then start it. Starting the server automatically creates a Default Hive with default paths and a Native-II vault with its other files and folders. Although this method is easy, it has the disadvantage of placing the initial Native-II hive on the same drive as the Native-1 vault, which may have many files and limited space.

Thus, the Hive Manager serves two main purposes

- Creating new hives for an existing 2005 or later server configuration, to increase the amount of available space.
- Creating an initial custom hive for a server configuration created with StarTeam 6.0 or earlier.

If you are using the Hive Manager to create a new hive, the server for which it is created must be running.

To use the *Hive Manager* to create a new hive:

1. Select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server. The *Server Administration* dialog appears.
2. Select the server configuration. If the configuration was created with StarTeam 6.0 or earlier, verify that it is stopped. StarTeam 2005 or later server configurations may be either running or stopped.
3. To display the *Hive Manager* dialog, click the Hive Manager shortcut in the lower left pane or select Tools > Administration > Hive Manager from the menu.button.
4. On the *Hive Manager* dialog, click New. This action displays the *New Hive* dialog. Note that the location of the hive-index.html file, which contains the properties for each hive used by the server configuration, appears near the top of the screen.
5. Enter information about the new hive in the following fields:

   Name: Unique name for the hive. DefaultHive is the default.

   Archive path: Path to the hive's Archives directory. The default path is <repository_path\DefaultHive\Archives.

   Cache path: Path to the hive's Cache directory. The default path is <repository_path\DefaultHive\Cache.

   Maximum cache size: Maximum number of megabytes of hard disk space that the Cache can use. The default is 20% of the disk space available. In the *Server Administration* dialog, you can calculate the correct default maximum size for the cache. However, if you are using the *Server Administration* dialog and it is not running on the same computer as the server, you cannot calculate the maximum size. In this situation, enter 100 MB, which is the default size for a Native-I vault.

   Cache cleanup interval: Seconds between cache cleanup/refresh operations. The default value is 600. The range is 60 (1 minute) to 3153600 (1 year).

   Storage limit threshold: Percentage of total disk space allowed for hive. When this percentage has been reached, no more archives can be added to the hive. The default is 95% of total disk space.
6. Select or deselect Allow new archives. The default is selected. If no other hives exist for the server configuration, this check box must be selected.
7. (Optional) If you are using StarTeamMPX and the Cache Agent, set a Root Cache Agent archive path.
8. Click OK to confirm your choices. This action returns you to the Hive Manager.
9. Click Close to return to the Server Administration dialog.

However, if you prefer, you can create a default hive for use with a server configuration created with StarTeam 6.0 or earlier. The disadvantage of this method is that it does not permit you to put the hive on another drive or volume.

To create a default hive for use with a server configuration created with StarTeam 6.0 or earlier:

1. Upgrade the server configuration to StarTeam 2005 or a later release. For details, see the *StarTeam Installation Guide*.
2. Select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server. The *Server Administration* dialog appears.
3. Select the server configuration.
4. If the server is running, click the Shut Down Server icon or select Actions > Shut Down Server from the menu.
5. Re-select the server and click the Start Server icon or select Actions > Start Server from the menu. This action automatically creates the Default Hive as well as the Native-II vault and its other files and folders. If you wish to view the properties of the Default Hive, you can do so with the Hive Manager.

Viewing Hive Properties

Sometimes you may want to review the properties of a specific hive or change its settings. For example, you may want to move its Archives or Cache folder to an alternate location. In that situation, you must use the Hive Manager to display the hive's properties, then change them.

Changing the Archives Path

Changing the Archives path for a hive is generally done because of a serious problem, such as a drive failure. It must also be done with caution, or the results can be unexpected.

Warning **A new archives path will not take effect until the server has restarted.**

For example, assume that the server is running and that you clear the "Allow new archives check box on the *Hive Properties* dialog. Then you copy the archive files to their new location, change the Archive path, and reselect "Allow new archives." All of these actions have been done without stopping and restarting the server configuration. Therefore, the new path will not take effect, but the hive immediately starts to accept new archives. Although the new path appears in the hive properties, it is not the path loaded from the .xml file. The server continues to use the prior path entered for the path, because it has not been restarted; thus, you have not accomplished your objective.

However, either of the following two scenarios allows you to successfully change the path used by the archive files:

Scenario 1:

Shut down the server, copy the archives files to the new location, change the archive path on the Hive Properties dialog, and restart the server. With this method, there is no need to deselect the "Allow new archives" check box, because no files can be added or checked in while the server is down.

Scenario 2

Leave the server running. Clear the "Allow new archives" check box (if at least one other hive exists) on the *Hive Properties* dialog. The files that are added or checked in will be sent to the other hive. At an appropriate time, shut down the server, copy the archive files to the new location, change the archive path, select the "Allow new archives" check box on the *Hive Properties* dialog, and restart the server. No files will be missed.

Changing Hive Properties

Hive properties can be changed only for a server that is running.

To view hive properties and change their settings:

1. Select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server. The *Server Administration* dialog appears.
2. Select the server configuration. The properties of local servers that are running or stopped can be changed, although changes on servers that are stopped will not go into effect until the next startup. Remote servers can be accessed only if they are running, and the changes to their properties occur immediately.
3. Click the Hive Manager shortcut in the lower left pane or select Tools > Administration > Hive Manager These actions display the *Hive Manager* dialog.
4. Click Properties. This action displays the *Hive Properties* dialog.
5. Review and, if desired, change the information in this dialog, which includes:

   Name

   Archive path

   Cache path

   Maximum cache size

   Cache cleanup interval

   Storage limit threshold

   All of the fields, except Name, are editable. For the default and possible settings for Maximum cache size, Cache cleanup interval, and Storage limit threshold, see "Creating New Hives" on page 267.
6. Select or deselect Allow new archives. The default is selected. Since a server configuration must have at least one active hive, you cannot clear this check box for all hives.
7. (Optional) If you are using StarTeamMPX, set a Root Cache Agent archive path.
8. Click OK to confirm your choices and return to the Hive Manager. If the server is local and running, the changes are made immediately. If it is local and not running, the changes take place when the server configuration is next started. Remote servers can be accessed only if they are running, and the changes occur immediately.

Converting Native-I Vault Archive Files

After you upgrade a server configuration created with StarTeam 6.0 or earlier to StarTeam 2005 or a later release, it uses the Native-II vault for all new files put under application control. However, it continues to use the old Native-I vault for existing files.

At some point, you will want to convert your Native-I archive files to Native-II. The conversion can be performed as a background process from the *Server Administration* dialog or manually from the client. When you convert a Native-I archive file to Native-II, you can delete the Native-I file. If it is not deleted, it is moved to a new subfolder of the Native-I vault Archive folder named ConvertedArchives.

Important The Server sends e-mail notification to members of the System Managers group after the vault conversion process is finished. By default, no users are in the System Manager group, so be sure to add some users to this group before the conversion process completes.

Converting Files as a Background Process

Borland recommends that you do a backup operation before converting your files. However, the conversion process **must not** be running during a backup operation. If you have scheduled the conversion process to run at specific times, check your scheduling program to be sure that the conversion is not scheduled to run during the backup.

This operation can be performed only on a server that is shut down.

To convert Native-I files as a background process:

1. On the computer where the Server is installed, select Start > Programs > StarTeam > StarTeam Server x.x > StarTeam Server. The *Server Administration* dialog appears.
2. Select a server configuration.
3. Select the Start Conversion icon or select Actions > Start Conversion. These items are enabled only for server configurations that contain one or more Native-I archive files.

   These actions display the *Start Conversion* dialog. They also toggle the Start Conversion button to Stop Conversion so that you can stop the conversion at any time.
4. On the *Start Conversion* dialog, select the Delete old archives after successfully converting them checkbox if you wish to retain a copy of the files in Native-I format.
5. Click OK to start the conversion process, which will run in the background. You can run the process until all of the files have been converted or stop it, at your convenience.

   Note: If the application determines that the server configuration does not yet have at least one hive, it will ask you to create one with the Hive Manager before converting the files.
6. To stop the conversion process (for example, during periods of high server activity), return to the Server Administration dialog and click Stop Conversion. You can restart the conversion again when it is convenient.

Using the ConvertVault Command-Line Application

A sample SDK application named ConvertVault comes with the Server. ConvertVault is a command-line application written in Java and stored in the Server installation folder. This application can start or stop the background vault conversion process. It accepts command-line parameters that allow you to specify the server address, user name and password, and so on. To learn more about how to use this application, type

```
ConvertVault -?
```

at the command prompt.

Because most Windows operating systems provide a scheduling application, you can schedule ConvertVault to start the conversion process when the server is least heavily used, such as overnight or on a weekend, and stop it again when appropriate.

As a precaution, do a complete backup before starting the background conversion process.

Running the Sample Application

In addition to the *.java source file, a compiled *.class file has been included. To run the application, you must have a compatible Java Runtime Environment (JRE) installed on your computer. For your convenience, a *.bat file has been included that allows you to

run the sample on Win32 platforms. You may need to edit the *.bat file to make it run properly on your computer.

If you are using the included *.bat file to run the sample application, be sure that the SDK_DIR variable points to the correct StarTeam SDK installation directory, that StJava.exe is properly installed in the SDK's bin folder, and that the bin folder is included in the PATH.

Disclaimer

This application is intended to provide an example of what can be accomplished using the StarTeam SDK. Borland does not guarantee the operability of the ConvertVault application and does not provide technical support for it. Borland is not responsible for data loss or damage caused by use of this sample application.

Converting Files Manually

Borland recommends that you always backup files before converting them. However, you should never convert files during a backup operation.

To convert selected files manually from the client:

1. Select the files you wish to convert.
2. Select File > Advanced > Convert. This action displays the Convert Archives dialog.
3. Select one of the following options:
   - Keep old archives
   - Delete old archives after successful conversion
4. Click Convert.

Note: To convert files manually, you must have access rights that allow you to perform remote configuration administration.

Verifying Tip Revision Information (Native-I Vaults)

When you use a Native-I vault and verify the files for a server configuration, Server Administration checks to see whether:

- The tip revisions of the files are in the vault.
- The database and vault use the same number for the tip revision of the file.
- The MD5 value for the tip revision is the same in both the database and the vault.
- Every cache file is identifiable.

If the answer to any of these inquiries is no, use Server Administration to attempt to repair the problem. Before you begin, however, be sure to create a backup of the database.

- If the tip revision is missing from the vault but can be located in the cache, Server Administration uses the cache file to recreate the tip revision in the vault.

  If the tip revision is not in the cache, Server Administration asks the user to check in the file, which may be in a working folder.
- If the vault and database have different numbers of revisions:
  - If the vault has too many revisions, the excess revisions in the vault are dropped until the database and the vault have the same number.

- If the database contains an extra revision and the tip revision can be located in the cache, Server Administration uses the cache file to recreate the tip revision in the vault.

- If the MD5 values differ and the tip revision is not already in the cache, Server Administration checks that revision out to the cache. The MD5 value is recalculated for the cached file and the database value is changed to the new value.
- If a cache file is "orphaned," Server Administration deletes it from the cache.

This operation can be performed only on servers that are not running.

To verify and repair tip revision information:

1. From the computer on which the Server is installed, go to the *Server Administration* dialog. See "Creating a Server Configuration" on page 12 for details.
2. From the list in the *Server Administration* dialog, select a server configuration to be verified.
3. Click the Verify Native-I Vault icon or select Actions > Verify Native-I Vault. The *Verification Status* dialog appears. Then a message indicates whether problems were found. If problems were found, the *Results* dialog appears.
4. Do one of the following:
   - Click Repair All to attempt to repair all the problems.
   - Select one or more problems to repair, then click Repair.

The Server Administration utility then fixes the problems, if possible, and places the word Fixed in the problem's Status column in the *Results* dialog.

If a file is corrupted beyond recognition, Server Administration cannot salvage it. If the file is not in the archive or the cache folders and you cannot specify a replacement, it remains missing. When a problem cannot be fixed, the Status column does not change.

Server Administration also creates a log for each session. The path to the latest log file for a server configuration is <repository path>\Logs\Repairs\Repair.log. If you have used Server Administration before to repair the server configuration, earlier log files contain the times of these repairs.

Purging Deleted Views from a Server Configuration

The Purge button on the Server Administration utility enables application administrators to remove deleted views from a server configuration database and vault and rebuild the indexes in that database. If the deleted view has items that are active in another view, these items are not deleted. For example, if two views share a file and you delete one view, the shared file is not deleted.

Borland recommends that you perform a purge after deleting one or more views from a project. The purge operation can take several hours if a large number of records need to be deleted or moved.

Be sure to backup the database before using the Purge feature. You should also start the server configuration from which the view was deleted at least once before using Purge.

Important Purge is available for Oracle, DB2, and Microsoft SQL Server databases. You must have installed the database client application on the same computer as the Server for Purge to work properly.

To purge deleted views from a database:

1. From the computer on which the Server is installed, go to the *Server Administration* dialog. See "Creating a Server Configuration" on page 12 for details.

2 If the server configuration is running, select it from the list box in the *Server Administration* dialog, and click the Shut Down Server icon or select Actions > Shut Down Server.

3 Re-select the server configuration you want to purge.

4 Click the Purge icon or select Actions > Purge. When requested, enter your user name and password.

Purging from the Command Line

A purge utility provided with StarTeam 2005 and later releases allows you to purge a server configuration from the command line. Typically, you place the command in a batch (.bat) file and schedule the batch file to run at a time when the server configuration will be shut down.

The syntax is:

StarTeamPurge *server_configuration_name*

The purge command must be run from the Server's installation folder. For example, if you wanted to purge the StarDraw sample server configuration, you would use:

```
StarTeamPurge StarDraw
```

or:

```
C:\ProgramFiles\Borland\StarTeamServer2005R2\StarTeamPurgeStarDraw
```

The second example assumes that the Server was installed at its default location.

Migrating a Server Configuration to Another Database

The Migrate button on the Server Administration utility enables you to migrate a server configuration from one database to another or from one database type to another type. The Migration button adds information about the new server configuration to the starteam-server-configs.xml file.

This operation can be performed only if the server is not running.

Before the migration:

1 Create a backup of the database you plan to migrate. Also ensure that you have backups of the files and folders in the server configuration repository. See Appendix C, “Backing Up the Application,” on page 249 for more information.

2 Ensure that the database to be migrated is still in good repair by doing one or both of the following:

   a On the *Server Administration* dialog, click Verify Native-I Vault. This button applies only to vaults created before StarTeam 2005 that have not yet been converted to Native-II architecture.

   b Run any verification tools provided by your database vendor.

3 Do one of the following:

   - Manually create a database or schema user as the recipient of the migrated data. See the *StarTeam Installation Guide* for details. Make sure that you note the names provided for the following:
     - DSN name
     - user name and password (for database or schema user)

       At a minimum, this user must have permission to create tables and stored procedures (if the database supports stored procedures).

- Use the Server to automatically create a database or schema user as the recipient of the migrated data.

To migrate a server configuration:

1. Plan the migration for a time at which it will inconvenience few users. A server configuration cannot be running while it is being migrated.

   Advise team members ahead of time that you plan to make the transition, let them know the time at which it will take place, and request that they check in their files.

2. Select Start > Programs > StarTeam > StarTeam Server *x.x* > StarTeam Server. The *Server Administration* dialog appears.

3. From the list box, select a server configuration to be migrated.

4. Click Migrate. A message warns you that you cannot migrate a server configuration if the server is not registered.

5. Click Yes if your server is registered.

   If your server is not registered:

   a. Click No.

   b. Click About.

   c. From the resulting *About StarTeam Server Administration* dialog, select the License node from the left pane.

   d. Click Register.

   e. In the resulting Server Registration dialog, enter your serial number and access key.

   f. Click OK and Close to return to the *Server Administration* dialog.

   g. Repeat step 3.

   The *Create a New Target Configuration for Migration* wizard appears.

6. In the *Select Target Configuration for Migration* dialog:

   a. Enter the name for the new server configuration in the Target Configuration Name text box.

   b. Click Next. The *Enter New Configuration Data* dialog appears.

7. Indicate the type of database and whether you want the application to create it for you.

   a. In the Database Type drop-down list box, select Microsoft SQL Server/MSDE, DB2, or Oracle.

   b. Do one of the following:

      - (Default and recommended action) Select the Create New StarTeam Database and ODBC Data Source check box.
      - Clear the check box if you have already manually created a database or schema user for this migration.

8. When this information is complete, click Next. From this point on, the dialog boxes are the same as those that appear when you create a server configuration. See the *StarTeam Installation Guide* for details about server configurations that use Microsoft SQL Server/MSDE, DB2, or Oracle.

When the migration is complete:

1. Disable the prior server configuration. This action prevents the server configuration from being started and accessed accidentally.

Important Both the old and the new server configurations access the same vault, cache, and attachments folders. However, they *do not* access the same database.

Continuing to use the prior server configuration will lead to vault verification errors and must be avoided!

a. On the *Server Administration* dialog, select the prior server configuration.

b. Click Disable/Enable. The Status column changes to “Disabled.”

2. Empty the Cache folder for the hive before starting the new server configuration. By default, the Cache folder is a child folder in the hive, under the repository root folder.

3. Run the following database script on the migrated database:

```
starteam_sqlserver_tuning_script
```

For information about what the script does, see the *StarTeam Installation Guide*.

4. Do **not** delete the repository, as the new server configuration uses the same repository as the original server configuration. Only the database changes.

5. After verifying that the new configuration works correctly, delete the prior server configuration, the database that it used, and its System DSN.

The Z99 table is a temporary table that records the progress of the database migration. If the Migrate process stops before completing, it uses the Z99 table to determine the point at which it should resume the migration when you restart the process.

Review the following columns to determine how far Migrate has progressed.

- Column 1 contains the source table name.
- Column 2 contains the ID of the last record copies.
- Column 4 contains either a Y or N, indicating whether the table copy is complete.

Exporting Database Information

The Catalog Export utility exports two application tables, Catalog_Tables and Catalog_Fields, into comma-delimited files. This tool is very useful for database administrators because Catalog Export translates database tables and column names into identifiers that are used by the Server.

If you examine a column of data in the exported field catalog and find that one record has a surprising value (for example, all other records have a -1 in a column, but this record has a 16-digit number), the record may have been corrupted. However, Borland does not recommend that you delete any records, especially if you have not backed up the database.

Exported data can be imported and viewed in any application that supports comma-delimited fields. For example, if you save the file with a .csv extension, it will open in Microsoft Excel.

This operation can be performed only if the server is not running.

To run Catalog Export:

1. From the computer on which the Server is installed, start Server Administration. See “Starting a Server Configuration” on page 15 for details.

2. From the list box in the *Server Administration* dialog, select a server configuration to be exported. The server must be shut down to perform this operation.

3. Click the Catalog Export icon or select Actions > Catalog Export . The *Catalog Export* dialog appears.

4. In the File Name For Exported Table Catalog text box, enter or browse for the location and path to which you want to export the table catalog.

5. In the File Name For Exported Field Catalog text box, enter or browse to the location and path to which you want to export the field catalog.

Note Be sure to enter the appropriate file extension for the application to which you want to export the files.

6 Click OK. The system displays the following message:

```
Catalogs exported successfully.
```

7 Open and view the files in the application in which you exported the files.

15

Using Import/Export Manager

StarTeam Import/Export Manager is a set of utilities that allow you to copy data from one StarTeam Server to another as a one-time necessity or as part of a continuous set of operations in which they keep the target server synchronized with the source server.

You must use a 1.4.x release of the Java VM with Import/Export Manager. Import/Export Manager can export from Servers that are release 5.2 and later. The import must be to a Server whose release is the same or later than that of the Server used for the export.

Flexible Extraction

You can extract the data from an entire server configuration, an individual project or view, or a branch of folders in a view. This extraction flexibility makes it possible for Import/Export Manager to support diverse use cases.

For example, you may copy a project from one server to another for archival purposes. In another situation, you might copy a project because the source server is overloaded and the users of the copied project will continue their work on another server.

You may copy a few folders to a target folder and do incremental updates to keep the files in those folders on the target server synchronized with the files on the source server. This is similar to sharing files between projects or views on one server.

In another situation, you might copy a few folders so that the target server has both the current files and their revision histories. From the target server, those files may be revised by new users to fit new applications.

Server Inaccessibility

When connections between users and the primary server cannot be maintained or are not allowed, Import/Export Manager can come to the rescue.

With Import/Export Manager, data export, transfer, and import are separate processes. Consequently, the source and target servers do not have to be accessible at the same time or from a single application. Data transfer can use a variety of techniques including file copy, FTP, or removable media.

One customer copies source code to several target servers around the world. Work continues on the next release on the source server. On the target servers, new child

views are created, and developers adjust the code to allow for regional differences, such as languages and state regulations. The product is then rebuilt and used regionally.

Another customer delivers source code to a military agency and, for security reasons, no connection between the customer server and the military server is allowed. With Import/Export Manager, the extraction of data from the source server results in a set of files that can be securely transported to the target server's location. Then, as a separate operation, Import/Export Manager can be used to import that data to the target server.

Installing Import/Export Manager

It is best to run the export and import utilities on the same computer as the Server that they are copying from or to. This is not a requirement, but running either utility on a different computer can degrade performance.

If you must run the utility on another computer, select one that is connected to the server via a high-speed LAN connection.

Install the same version of Import/Export Manager for use with both the source and target servers.

Importing and Exporting

Import/Export Manager utilities allow you to "export" an existing StarTeam source object (server configuration, project, view or folder) and later import that data into a StarTeam target object (server configuration, view or folder). Over time, you can synchronize the source and target objects by using delta export and import operations that apply changes made to the source object to the target object. The changes are those made to the source object since the initial full export operation or the last delta export operation. All changes to the target object must be made by the import utility. No manual check-in operations or item creation should occur in the target object.

The export utility allows you to specify a source object and write its contents in a canonical format. This canonical format is represented by a series of XML documents and other files.

The import utility allows you to take the output from the export utility and import it into a server configuration, view or folder (the target object), creating a copy of the original source object. Be aware that after running the import utility, the target object may not be an exact duplicate of the source object. See "Understanding Import/Export Manager Limitations" on page 286.

Delta operations from a source object to a target object are the same as full operations with the following exceptions:

- In a delta export operation, only the changes to the source object that occurred since the last export of the source object are recorded.
- In a delta import operation, only the changes recorded by the corresponding delta export are imported into the target object.

This section provides an overview of general Import/Export Manager process and how to use the export and import utilities. Subsequent sections describe in greater detail how to export and import a server, project, view or folder.

Using the Export Utility

The export utility copies most of the data from the source object to a set of XML and other files. For example, contents of file revisions are preserved in individual files.

You do not have to lock the server configuration before running the export utility. However, if users are still working, be aware of the following:

- The time at which the export utility starts determines what objects and items it copies. For example, if a view is created after the export utility starts, that view's contents are not copied.
- Some objects, such as access rights, have no time information associated with them. The export utility copies this data "as-is" at the time the utility makes the call to the server. The information may have changed between the time the utility started and the time the call was made.

In the following procedures, the Import/Export Manager's installation folder is referred to as your STIEM_HOME directory. For best results your STIEM_HOME directory should be on the computer that is running the server configuration you wish to export.

To use the export utility:

1. Using your favorite text editor, create a "export.cdl" file for the type of export (server, project, view or folder). The next few sections describe the contents of the export.cdl file for the specific type of export you wish to perform.

Important Make sure you have adequate disk space when running the export utility. Make sure your output directory is located on a hard drive that has enough available capacity to store all of the output. A conservative estimate of the available space you will need is twice the size of your server's Archive and Attachments folders.

2. Open a command prompt and change to your STIEM_HOME directory.
3. At the command prompt, type the following command:
   ```
   stexport export.cdl
   ```
4. When the export utility completes, you can see that it generated output in the directory specified in your export.cdl file.

Copying the Data

Once the export utility completes, you must copy the contents of the output directory to the target server. If you need to copy the data over a slow network connection or if there is a large amount of data, consider compressing the output folder before copying the data.

Using the Import Utility

The import utility copies most of the data from the set of XML and other files created by the export utility to the target object specified in the import.cdl file you will create for the type of import. Please be aware that depending on the type of import you are running (server, project, view or folder) there are certain restrictions placed on the target object. The next few sections describe the contents of the import.cdl file for the specific type of export you wish to perform.

In the following procedures, Import/Export Manager's installation folder is referred to as your STIEM_HOME directory. For best results your STIEM_HOME directory should be on the computer that is running the server configuration on which the data will be imported.

Caution You cannot run import operations against the same target object using different source objects.

To run the import utility:

1. Using your favorite text editor create an import.cdl file in your STIEM_HOME directory called for the type of import you are doing (server, project, view or folder). See the other sections for a description of the contents of the import.cdl file for the specific type of import you wish to perform.
2. Open a command prompt and change to your STIEM_HOME directory.
3. At the command prompt, type the following command:
   ```
   stimport import.cdl
   ```
4. When the import utility completes, you can connect to your target server with any client and see that contents from the source server have been copied into the target object specified in your import.cdl file.

Creating the .cdl Files

The export.cdl and import.cdl files vary depending on the scope of the operations. Whether you are exporting and importing a server configuration, project, view, or folder determines what parameters you use. The next few sections provide step-by-step directions for each scope.

Some parameters, such as EXPORT_RECORD_FILE and IMPORT_RECORD_FILE are used in only the export.cdl or only the import.cdl file, but most parameters apply to either file.

With the exception of the parameters that define the scope of the export or import operation (PROJECT, VIEW, and FOLDER) which you use as needed, all the parameters will be used in at least one of the .cdl files.

The order of the parameters in the .cdl file is not important, but each parameter appears on a separate line of the file. Each parameter uses the following syntax:

parameterName=*value*

The parameter and its value are separated by the equals symbol. No delimiters, such as quotation marks, are required for any string values.

The following table lists the parameters alphabetically and explains their purpose.

Table 15.1 Parameters for all CDL Files

Parameter	Purpose
CR_MAPPING_FIELD	See “Mapping Items” on page 284.
CR_MAPPING_ PLACEMENT	See “Mapping Items” on page 284.
CR_MAPPING_STRING	See “Mapping Items” on page 284.
EXPORT_RECORD_ FILE	Used only in export.cdl files. Set to the path to a record file that the export utility uses to store the times at which it runs an export against a particular source object. The utility creates this file for you during a full export operation; you simply provide the name. You specify the same file in each subsequent delta export operation. Never manually edit the file specified as either the export or import record file. If you run a full export operation against a source object and then need to run another full export operation against the same source object in the future, the second time you must specify a different EXPORT_RECORD_FILE than you did for the first full export operation.

Table 15.1 Parameters for all CDL Files (continued)

Parameter	Purpose
EXTRACT_DISABLED_ PROPERTIES	Used only in export.cdl files. Indicates how to handle disabled custom fields. When set to true, the default, the export utility records all custom fields and they show up on the target server as disabled custom fields after your run the import utility. When set to false, the export utility does not record disabled custom fields. The target server has no knowledge of them. The setting for this feature must be the same in the initial export and in all delta exports to the same target server.
FOLDER	Used only when exporting or importing a folder. Set to the path to the folder in the hierarchy. For example, suppose you have a project named “Test” and you want to export or import its root folder. Assuming the root view is named “Test” and the root folder is named “Test”, the three lines in your export.cdl file would look be: PROJECT=Test VIEW=Test FOLDER=Test
INCLUDE_USER_ DATA	Meaningful only when the export or import involves a subset of a server configuration. When you export or import an entire server configuration, the user data is always included. When this option is used to export or import a part of the server configuration, its existence is important, not its value. For example, you could set the option to “1” or “xyz”. Regardless of the value, information about users is added to the export output or to the target server configuration’s repository. The export utility’s output includes all user properties (except for passwords), all group properties, and all access rights. During the import operation, each user’s password is set to the user name. During the import operation, item properties that contain user information, such a the Modified By and Created By properties, are set to the user value as in the source. When this option does not appear in the .cdl file, user information is ignored. Item properties that contain user information are set to the user name of the user running the import utility.
INPUT_DIRECTORY	Used only in import.cdl files. Set to the path to the directory to which you copied the export output.
IMPORT_RECORD_ FILE	Used only in import.cdl files. Set to the path to a record file that the import utility uses to map IDs between the source and target objects. The utility creates this file for you during a full import operation; you simply provide the name. You specify the same file in each subsequent delta import operation. Never manually edit the file specified as either the export or import record file.
LOG_PROPERTIES	Set to the full path to the import.log.properties file that is located in your installation folder. For more information about this file, see “Setting the Log Properties” on page 286.
MODE	Set to either FULL or DELTA. Use FULL only when you are first running an export or import operation against a given source object. Use DELTA only when you have run at least one previous export or import operation against the source object specified in the configuration file.

Table 15.1 Parameters for all CDL Files (continued)

Parameter	Purpose
OUTPUT_DIRECTORY	Used only in export.cdl files. Set to the path to the directory where you want to have the export output written. If that directory does not exist, the export utility will create it for you.
PASSWORD	Set to an Administrator's password.
PROJECT	Used when exporting or importing a project, view, or folder. Set to the name of the project.
REQ_MAPPING_FIELD	See "Mapping Items" on page 284.
REQ_MAPPING_ PLACEMENT	See "Mapping Items" on page 284.
REQ_MAPPING_ STRING	See "Mapping Items" on page 284.
SERVER_NAME	Set to the name of the server. If you installed the export utility on the same computer as your server, you can use the string "127.0.0.1".
SERVER_PORT	Set to the port number of the server. The default server configuration port is 49201, although this value may be different on your computer.
TASK_MAPPING_ FIELD	See "Mapping Items" on page 284.
TASK_MAPPING_ PLACEMENT	See "Mapping Items" on page 284.
TASK_MAPPING_ STRING	See "Mapping Items" on page 284.
USERNAME	Set to an Administrator's user name.

VIEW	Used when exporting or importing a view or folder.
	Set to the path from the project's root view to the view to be exported or imported. This string's format should be `<root view>/<child view>/<child view>/....`
	For example, suppose you have a project name "My Project" with its root view named "My Project" and a child view called "ChildView 1". To export or import ChildView 1, use:
	`PROJECT=My Project`
	`VIEW=My Project/ChildView 1`

Mapping Items

Import/Export Manager cannot guarantee that numbered items, such as change requests, will have the same numbers on both the source and target servers. In fact, it is very likely that they will not. Because of this the import.cdl file can include mapping parameters for each type of item that is numbered: change requests, requirements, tasks, and topics.

Using the mapping parameters, you specify a text field for the item on the target server as the field that will contain information about the item's number on the source server. For example, you might specify the "External reference" field of a change request.

In addition, you specify a string value for that field. It can be any string value that will fit in that field without truncation, but it must contain appropriate terms that for the type of item and the source server. The import utility substitutes the number of the item and the address and port of the source server for these terms.

For example, <CR> represents the number of the change request, and <SOURCE> represents the address and port number of the source server. If you use the following string:

```
Was CR-<CR> on source server <SOURCE>
```

the import utility will place a value such as the following in the specified field:

```
Was CR-1587 on source server 127.0.0.1:49201
```

However, it is possible that the specified field already contains data. If so that is data that you probably do not want to lose. So Import/Export Manager provides another parameter to indicate whether you want the string prepended or appended to the existing value in that field. You can also specify that the existing data will be overwritten.

When you use the parameters in the following table, you substitute *item* for CR, REQ, TASK, or TOPIC.

Table 15.2 Mapping Parameters for import.cdl File

Parameter	Purpose
*item*_MAPPING_FIELD	Specifies the internal field name of the field that, on the target server, will be used for mapping purposes. The field must be a valid text field for the item type. For example, a custom field may be created for the item type on the target server.
*item*_MAPPING_STRING	Specifies the string to be used in the mapping field. Must not exceed the length of the field. It must contain a term indicating the item type that matches the item type in the parameter name: <CR>, <REQ>, <TASK>, or <TOPIC>. For example, CR_MAPPING_STRING must contain <CR>. The import utility substitutes the number of the item for this term. The string may optionally contain the term <SOURCE> which will be replaced with the source server's address and port number, separated by a colon.
*item*_MAPPING_ PLACEMENT	Set to PREPEND, APPEND, or OVERWRITE to indicate where the mapping string will appear in the mapping field. Use PREPEND to place the mapping string in front of data already in the mapping field. Use APPEND to place the string after data already in the mapping field. Use OVERWRITE to remove the existing data in the mapping field. The mapping string will be the field's entire contents.

CR_MAPPING_FIELD, CR_MAPPING_STRING, and CR_MAPPING_PLACEMENT are the fields for change requests and, when used, all three parameters must be included in the import.cdl file. CR_MAPPING_STRING must contain the term <CR>.

REQ_MAPPING_FIELD, REQ_MAPPING_STRING, and REQ_MAPPING_PLACEMENT are the fields for requirements and, when used, all three parameters must be included in the import.cdl file. REQ_MAPPING_STRING must contain the term <REQ>.

TASK_MAPPING_FIELD, TASK_MAPPING_STRING, and TASK_MAPPING_PLACEMENT are the fields for tasks and, when used, all three parameters must be included in the import.cdl file. TASK_MAPPING_STRING must contain the term <TASK>.

TOPIC_MAPPING_FIELD, TOPIC_MAPPING_STRING, and TOPIC_MAPPING_PLACEMENT are the fields for topics and, when used, all three parameters must be included in the import.cdl file. TOPIC_MAPPING_STRING must contain the term <TOPIC>.

Setting the Log Properties

The Import/Export Manager utilities use the Apache Software Foundation's log4j logging libraries. The logging is controlled by a log4j properties file. The utilities ship with two sample log4j properties files:

- export.log.properties
- import.log.properties

To learn how to modify these files to meet your needs please visit: http://jakarta.apache.org/log4j/docs/index.html and check the online documentation.

As a quick example, suppose you wanted to suppress entirely the output to the command-line from the export utility. To do so, simply change the following line the export.log.properties file from:

```
log4j.rootLogger=INFO, stdout, file
```

to:

```
log4j.rootLogger=INFO, file
```

Understanding Import/Export Manager Limitations

The utilities cannot accurately copy all server information because of deletions, moved items, and shared items.

Deletions

Import/Export Manager has several limitations when it comes to handling deleted data.

- Deleted data and view configurations

  The utilities do not track deleted data. For example, a file item that was deleted on the source server can still be found by rolling back the view configuration to a time before the item existed. Rolling back the view configuration on the target server would not reveal the deleted file.

  For example, suppose you created a file at 8 a.m. on Jan. 1, 2003 on the source server. Now suppose you deleted that file at 10 a.m. on Jan. 1, 2003. Rolling the view configuration on the source server back to Jan. 1, 2003 at 9 a.m. will allow you to access that file.

  If you ran an export and import of that source server to a target server sometime after 10 a.m. on Jan. 1, 2003 and then rolled back the view configuration on the target to Jan. 1, 2003 9 a.m., you do not find that file on the target server.

- Deleted data and object IDs

  Deleted objects (such as user accounts, group accounts, projects, views, and items) on the source server skew the object ID allocations for those objects on the target server. Below is a simple scenario to show how IDs can become out of synchronization.

  Suppose the source server has a single project with a single view and a single folder, the root folder. The root folder has object ID 0, and item ID 0.

  Suppose three change requests are created in order: A, B and C. Change Request A has object ID 1, item ID 1, and change request number 1. Change Request B has object ID 2, item ID 2, and change request number 2. Change Request C has object ID 3, item ID 3, and change request number 3.

  Now suppose that Change Request B is deleted.

  After running the utilities, the target server configuration looks like this:

The root folder has object ID 0, and item ID 0.

Change Request A has object ID 1, item ID 1, and change request number 1. Change Request B does not get copied. Change Request C has object ID 2, item ID 2, and change request number 2.

- Deleted user IDs

  Any reference to a deleted user ID on the source server is replaced by the user ID used by the import utility to connect to the target server.

  For example, assume user John Doe (User ID 10) creates change request 123 on the source server and then his user account is deleted. When change request 123 gets replicated to the target server, change request 123 looks like it was created by an existing user, the one whose user ID was used to authenticate the import utility.

  Customers may want to use a special ID when running an import to distinguish between IDs that were substituted for deleted IDs during the import and data that was routinely entered by the user who just happens to have also run the import utility.

- Deleted items and reference trees

  On the source server, suppose you create a file, File F, in Folder A. Then you share File F into Folder B and delete File F in Folder A.

  At this point, if you look at the item reference tree for File F in Folder B, you see that it has been shared from the deleted File F in Folder A.

  After an export and import, the tree of File F in Folder B looks as though File F has always been in Folder B. File F in Folder B is the root of its Reference Tree in the target server.

Moved Items

An item that was moved from Folder A to Folder B on the source looks like it has always existed in Folder B on the target.

For example, suppose you add a new file, File F, to Folder A at 8 a.m. On the source server. At 9 a.m., you create a new folder, Folder B, and move File F to Folder B.

The history for moved items does not look the same when you roll back the view configuration on both the source and target servers to a time at which File F was in the Folder A. On the target server, File F appears in Folder B. It is as if the file were never in Folder A.

Shared Items

The first limitation regarding shared items is explained in “Deletions” on page 286 because it also involves a deleted container.

The other limitation is that the exact time of the share is not recorded. The only way to tell that the share times of an item on the source and the target are not equal is to rollback the view configurations on the source and target views to the time when the source item was shared and see that the shared item exists on the source but not on the target.

Illegal XML Characters

XML 1.0 does not support all Unicode characters. From the XML 1.0 Specification:

```
Char ::= #x9 | #xA | #xD | [#x20-#xD7FF] | [#xE000-#xFFFD] | [#x10000-#x10FFFF]
```

Characters outside of this range can be included in item property fields, although in practice this is unlikely to occur. In the case that such characters do exist, however, the

export and import utilities cannot include them in the XML-based canonical representation.

The approach the utilities take is to replace these illegal XML characters with the character “?”.

Known Issues

The export and import utilities will fail if they are run against an “invalid” server configuration. An invalid server configuration has an invalid database or a repository that is either corrupted or out of synchronization with the database.

If you run the export or import utility on a different computer than Server, make sure that the system times on both computers are synchronized. Otherwise, you may experience errors while running the utilities.

Using Server Scope

This section describes the steps necessary to export and import the contents of one server configuration.

To create the export.cdl file:

1. Using your favorite text editor, create a file in your STIEM_HOME directory called export.cdl. This file is the export utility command file.
2. Copy and paste the following text into the export.cdl file:

```
OUTPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
EXPORT_RECORD_FILE=<export record file>
MODE=<mode>
EXTRACT_DISABLED_PROPERTIES=<TF>
```

3. In the export.cdl file, replace <output directory> with the path to the directory where you would like to have the export output written. If that directory does not exist, the export utility will create it for you.
4. In the export.cdl file, replace <server name> and <server port> with the name and port number of the server. If you installed the export utility on the same computer as your server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.
5. In the export.cdl file, replace <username> and <password> with the user name and password of an Administrator for the server. The default user name and password are “Administrator” and “Administrator”.
6. In the export.cdl file, replace <log properties file> with the full path to the export.log.properties file that is located in your STIEM_HOME directory.

   For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/export.log.properties”.
7. Replace <export record file> with the path to your export record file. See “Creating the .cdl Files” on page 282 for details about this file.
8. Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.

9 Replace <TF> with true to export information about all custom fields or false to exclude information about disabled custom fields. This setting must be the same for the initial and all the delta exports between the source and target servers.

10 Save the export.cdl properties file.

To create the import.cdl file:

1 Using your favorite text editor create a file in your STIEM_HOME directory called import.cdl. This file is your import utility command file.

2 Copy and paste the following text into the import.cdl file:

```
INPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
IMPORT_RECORD_FILE=<import record file>
MODE=<mode>
```

3 In the import.cdl file, replace <input directory> with the path to the directory to which you copied the export output.

4 In the import.cdl file, replace <server name> and <server port> with the name and port number of your target server.

If you installed the import utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.

5 In the import.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.

6 In the import.cdl file, replace <log properties file> with the full path to the import.log.properties file that is located in your STIEM_HOME directory.

For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/import.log.properties”.

7 Replace <import record file> with the path to your import record file. See “Creating the .cdl Files” on page 282 for details about this file.

8 Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.

9 Save the import.cdl properties file.

Using Project Scope

This section describes the steps necessary to export and import the contents of one StarTeam project.

To create the export.cdl file:

1 Using your favorite text editor create a file in your STIEM_HOME directory called export.cdl. This file is your export utility command file.

2 Copy and paste the following text into the export.cdl file:

```
OUTPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
```

```
PROJECT=<project name>
EXPORT_RECORD_FILE=<export record file>
MODE=<mode>
INCLUDE_USER_DATA=1
EXTRACT_DISABLED_PROPERTIES=<TF>
```

3 In the export.cdl file, replace <input directory> with the path to the directory to which you copied the export output.

4 In the export.cdl file, replace <server name> and <server port> with the name and port number of your target server.

   If you installed the export utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.

5 In the export.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.

6 In the export.cdl file, replace <log properties file> with the full path to the export.log.properties file that is located in your STIEM_HOME directory.

   For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/export.log.properties”.

7 In the export.cdl file, replace <project name> with the name of the project that you wish to export. If no such project exists, the export utility will report an error and exit.

8 Replace <export record file> with the path to your export record file. See “Creating the .cdl Files” on page 282 for details about this file.

9 Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.

10 If you do not want to include information about users, groups, and access rights, remove the INCLUDE_USER_DATA option.

11 Replace <TF> with true to export information about all custom fields or false to exclude information about disabled custom fields. This setting must be the same for the initial and all the delta exports between the source and target servers.

12 Save the export.cdl properties file.

To create the import.cdl file:

1 Using your favorite text editor create a file in your STIEM_HOME directory called import.cdl. This file is your import utility command file.

2 Copy and paste the following text into the import.cdl file:

```
INPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
IMPORT_RECORD_FILE=<import record file>
MODE=<mode>
INCLUDE_USER_DATA=1
```

   You do not use a PROJECT parameter because a project can be imported into a server configuration, but not into a parent project.

3 In the import.cdl file, replace <input directory> with the path to the directory to which you copied the export output.

4 In the import.cdl file, replace <server name> and <server port> with the name and port number of your target server.

If you installed the import utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.

5 In the import.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.

6 In the import.cdl file, replace <log properties file> with the full path to the import.log.properties file that is located in your STIEM_HOME directory.

For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/import.log.properties”.

7 Replace <import record file> with the path to your import record file. See “Creating the .cdl Files” on page 282 for details about this file.

8 Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.

9 If there is no user data in the export files, or if you do not want to import information about users, groups, and access rights, remove the INCLUDE_USER_DATA option.

10 Save the import.cdl properties file.

Additional Project Scope Export/Import Restrictions

In addition to the limitations mentioned in “Understanding Import/Export Manager Limitations” on page 286, the Project Scope has additional limitations.

- If your target server already has a project with same name as the project you are trying to import, an error message will be generated and the import utility will exit.
- Links between items in the exported project and items in other projects will not be copied.
- Items shared into the project you are exporting from another server configuration will not be treated as shares. They will be treated as if they were originally created in the project you are exporting.

Using View Scope

This section describes the steps necessary to export and import the contents of one StarTeam view.

To create the export.cdl file:

1 Using your favorite text editor create a file in your STIEM_HOME directory called export.cdl. This file is your export utility command file.

2 Copy and paste the following text into the export.cdl file:

```
OUTPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
PROJECT=<project name>
VIEW=<view path>
EXPORT_RECORD_FILE=<export record file>
MODE=<mode>
INCLUDE_USER_DATA=1
```

```
EXTRACT_DISABLED_PROPERTIES=<TF>
```

3. In the export.cdl file, replace <input directory> with the path to the directory to which you copied the export output.
4. In the export.cdl file, replace <server name> and <server port> with the name and port number of your target server.

   If you installed the export utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.
5. In the export.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.
6. In the export.cdl file, replace <log properties file> with the full path to the export.log.properties file that is located in your STIEM_HOME directory.

   For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/export.log.properties”.
7. Replace <project name> with the name of the project containing the view that you wish to export. If no such project exists, the export utility will report an error and exit.
8. Replace <view path> with the path from the root view to the view you wish to export. The <view path> string should be of the format <root view>/<child view>/<child view>/....

   For example, suppose you have a project name “My Project” with its root view named “My Project” and a child view called “ChildView 1".

   To export the view named ChildView 1, use:
   ```
   PROJECT=My Project
   VIEW=My Project/ChildView 1
   ```
9. Replace <export record file> with the path to your export record file. See “Creating the .cdl Files” on page 282 for details about this file.
10. Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.
11. If you do not want to include information about users, groups, and access rights, remove the INCLUDE_USER_DATA option.
12. Replace <TF> with true to export information about all custom fields or false to exclude information about disabled custom fields. This setting must be the same for the initial and all the delta exports between the source and target servers.
13. Save the export.cdl properties file.

To create the import.cdl file:

1. Using your favorite text editor create a file in your STIEM_HOME directory called import.cdl. This file is your import utility command file.
2. Copy and paste the following text into the import.cdl file:
   ```
   INPUT_DIRECTORY=<output directory>
   SERVER_NAME=<server name>
   SERVER_PORT=<server port>
   USERNAME=<username>
   PASSWORD=<password>
   LOG_PROPERTIES=<log properties file>
   PROJECT=<project name>
   VIEW=<view path>
   IMPORT_RECORD_FILE=<import record file>
   MODE=<mode>
   INCLUDE_USER_DATA=1
   ```

3. In the import.cdl file, replace <input directory> with the path to the directory to which you copied the export output.
4. In the import.cdl file, replace <server name> and <server port> with the name and port number of your target server.

   If you installed the import utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.
5. In the import.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.
6. In the import.cdl file, replace <log properties file> with the full path to the import.log.properties file that is located in your STIEM_HOME directory.

   For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/import.log.properties”.
7. Replace <Project Name> with the name of the project containing the target object view.
8. Replace <View Path> with the path from the target root view to the target object view on your target server where you wish the source object (the exported view) to be imported. This is the path to the view that will be the imported view’s parent.
9. Replace <import record file> with the path to your import record file. See “Creating the .cdl Files” on page 282 for details about this file.
10. Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.
11. If there is no user data in the export files, or if you do not want to import information about users, groups, and access rights, remove the INCLUDE_USER_DATA option.
12. Save the import.cdl properties file.

Additional View Scope Export/Import Restrictions

In addition to the limitations mentioned in “Understanding Import/Export Manager Limitations” on page 286, the View Scope has additional limitations.

- If your target view already has a child view with same name and path as the view you are trying to import, an error message will be generated and the import utility will exit.
- Links between items in the exported view and items in other views will not be copied.
- Items shared into the view you are exporting from another view will not be treated as shares. They will be treated as if they were originally created in the view you are exporting.
- When you import a view into a target view, it will always be created as a non-derived view.

Using Folder Scope

This section describes the steps necessary to export and import the contents of one StarTeam folder and its child folders.

To create the export.cdl file:

1. Using your favorite text editor create a file in your STIEM_HOME directory called export.cdl. This file is your export utility command file.
2. Copy and paste the following text into the export.cdl file:

```
OUTPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
PROJECT=<project name>
VIEW=<view path>
FOLDER=<folder path>
EXPORT_RECORD_FILE=<export record file>
MODE=<mode>
INCLUDE_USER_DATA=1
EXTRACT_DISABLED_PROPERTIES=<TF>
```

3. In the export.cdl file, replace <input directory> with the path to the directory to which you copied the export output.
4. In the export.cdl file, replace <server name> and <server port> with the name and port number of your target server.

   If you installed the export utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.
5. In the export.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.
6. In the export.cdl file, replace <log properties file> with the full path to the export.log.properties file that is located in your STIEM_HOME directory.

   For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/export.log.properties”.
7. Replace <project name> with the name of the project containing the folder that you wish to export. If no such project exists, the export utility will report an error and exit.
8. Replace <view path> with the path from the root view to the view containing folder you wish to export. The <view path> string should be of the format <root view>/<child view>/<child view>/....

   For example, suppose you have a project name “My Project” with its root view named “My Project” and a child view called “ChildView 1".

   To export a folder from the view named ChildView 1, use:

```
PROJECT=My Project
VIEW=My Project/ChildView 1
```

9. Replace <folder path> with the path from the root folder to the folder you wish to export.

   For example, suppose you have a project named “Test” and you want to export its root folder. Assuming the root view is named “Test” and the root folder is named “Test”, the three lines in your export.cdl file would like this:

```
PROJECT=Test
VIEW=Test
FOLDER=Test
```

10. Replace <export record file> with the path to your export record file. See “Creating the .cdl Files” on page 282 for details about this file.
11. Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.

12 If you do not want to include information about users, groups, and access rights, remove the INCLUDE_USER_DATA option.

13 Replace <TF> with true to export information about all custom fields or false to exclude information about disabled custom fields. This setting must be the same for the initial and all the delta exports between the source and target servers.

14 Save the export.cdl properties file.

To create the import.cdl file:

1 Using your favorite text editor create a file in your STIEM_HOME directory called import.cdl. This file is your import utility command file.

2 Copy and paste the following text into the import.cdl file:

```
INPUT_DIRECTORY=<output directory>
SERVER_NAME=<server name>
SERVER_PORT=<server port>
USERNAME=<username>
PASSWORD=<password>
LOG_PROPERTIES=<log properties file>
PROJECT=<project name>
VIEW=<view path>
FOLDER=<folder path>
IMPORT_RECORD_FILE=<import record file>
MODE=<mode>
INCLUDE_USER_DATA=1
```

3 In the import.cdl file, replace <input directory> with the path to the directory to which you copied the export output.

4 In the import.cdl file, replace <server name> and <server port> with the name and port number of your target server.

If you installed the import utility on the same computer as your target server, you can use the string “127.0.0.1” for the <server name> property. The default server configuration port is 49201, although this value may be different on your computer.

5 In the import.cdl file, replace <username> and <password> with the user name and password of an Administrator for your target server. The default user name and password are “Administrator” and “Administrator”.

6 In the import.cdl file, replace <log properties file> with the full path to the import.log.properties file that is located in your STIEM_HOME directory.

For example, if your STIEM_HOME directory is “C:\cdl”, replace the <log properties file> tag with the string “C:/cdl/import.log.properties”.

7 Replace <Project Name> with the name of the project that will contain the new folder.

8 Replace <View Path> with the path to from the target root view the view containing target object folder where you wish the source object (the exported folder) to be imported.

9 Replace <folder path> with the path from the root folder to the folder where you wish the source object (the exported folder) to be imported. This is the path to the folder that will be the imported folder’s parent.

10 Replace <import record file> with the path to your import record file. See “Creating the .cdl Files” on page 282 for details about this file.

11 Replace <mode> with either FULL or DELTA, depending on whether you are doing a complete or incremental export.

12 If there is no user data in the export files, or if you do not want to import information about users, groups, and access rights, remove the INCLUDE_USER_DATA option.

13 Save the import.cdl properties file.

Additional Folder Scope Export/Import Restrictions

- In addition to the limitations mentioned in “Understanding Import/Export Manager Limitations” on page 286, the Folder Scope has additional limitations.
- If there exists in your target folder a sub-folder with same name as the folder you are trying to import, an error message will be generated and the import utility will exit.
- Links between items in the exported folder and items not contained in the source folder hierarchy will not be copied.
- Items shared into the folder you are exporting from another folder will not be treated as shares. They will be treated as if they were originally created in the folder you are exporting.

Appendix

F

Entering Change Requests

The Change Request component makes it easy to track problems and correct them. However, you usually need a set of procedures that describe how your organization enters and uses change requests. Therefore, this chapter presents:

- A sample procedure for entering a change request.
- A table of the fields contained in the *Change Request* dialog.
- A syntax/shorthand table for entering defect and change request descriptions.

You can copy and distribute the sample procedure to users as is or modify it to reflect the preferences of your organization.

It may be necessary to expand the sample to explain additional tasks. For example, the workflow automatically makes the person who enters a change request responsible for verifying it. Your organization may require a different system, such as requiring that a particular user be responsible for change requests that have been resolved. In this situation, your set of procedures needs to include this information.

For a detailed discussion of the Change Request component, see "Tracking Change Requests," in the *StarTeam User's Guide*.

Sample Procedure: How to Enter a Change Request

As part of your job, you will enter change requests to record problems and suggested changes to the product on which your team is working.

This example illustrates how you might enter a change request for the following problem:

"When I choose which change request fields to display, the Available fields disappear if I check and uncheck the Show Advanced Fields box."

To enter a change request:

1. In a client, select a folder from the folder hierarchy.
2. Select the Change Request tab on the upper pane. The change request list appears in the upper pane.
3. Select New from the Change Request or context menu. The *New Change Request* dialog appears.

4 Fill in as many fields as you can or as many as your team requires. For a description of each field, see Table 15.3, “Change Request Fields,” on page 298.

5 Click Apply to save the changes, and then click OK to exit the dialog.

Using Change Request Fields

The following table lists the fields on the *Change Request* dialog, explains their uses, and indicates which fields are required by the application. In the table, CR is used as an acronym for change requests.

Table 15.3 Change Request Fields

Field	Req	Description	Example
Status	Y	For new CRs, set the Status field to New. The Status is changed to Open when the CR is assigned to a developer.	In this example, the status should be New.
Severity	Y	The Severity of a CR reflects the seriousness of the problem. High severity items are usually associated with data loss or corruption, system crashes, etc. Low severity items are generally misspelled items and cosmetic errors.	In this example, the problem is comparatively minor (that is, if it does not cause the system to crash or lose data), so classify it as Medium.
Priority	Y	In most defect tracking systems, Priority is a multi-level choice (usually on a 1 to 5 scale). In the application, however, it is a Yes or No choice. The priority of a CR is sometimes determined by the tester and sometimes by the developer. In most cases, it reflects the need to get a particular defect fixed before others. If the defect is catastrophic or prevents your team from accessing other major areas of the application, select the Priority field.	In this example, leave the Priority field cleared.
Platform	Y	The Platform indicates what type of operating system environment the defect occurs in. If the defect happens only on Windows NT, select Windows NT. In most cases, the defect will appear on all platforms.	In this example, set Platform to All.
External Reference	N	An External Reference refers to information received from outside the company, such as a note about a defect from an outside testing service or a customer. Currently this field is not used.	In this example, leave the field empty.
Component	N	The Component field identifies the component of the product in which the defect occurs. Currently this field is not used.	In this example, leave the field empty.
Category	N	The Category field identifies a subcomponent of the product. It is used with the Component field to identify the location in which the defect occurs. Currently this field is not used.	In this example, leave the field empty.
Synopsis	Y	The Synopsis field is used to give a *brief* summary of the problem encountered or the suggested enhancement. Consider the synopsis to be a title for the defect. **Note**: The Synopsis should only contain information for one defect. If the reported defect uncovers or relates to another defect, the second defect should be written up separately and referenced to the first defect in the synopsis (for example, “CR #3109 also relates to this defect”).	For this example, a synopsis might be: “Available fields disappear when using the Advanced Fields box.”

Table 15.3 Change Request Fields (continued)

Field	Req	Description	Example
Type	Y	If the CR is a reproducible problem in the software, select Defect. If it is a customer request or a feature enhancement request, select Suggestion.	For this example, select Defect.
Last Build Tested	Y	This field indicates the build number of the software in which the defect was discovered or last tested. If you are writing a CR, select the build number from the application (often found in the *About* dialog). If you are verifying or regressing the CR, and the problem still exists in the current build, change this field to the build number you are currently testing.	For this example, select the most current build number.
Addressed in Build	Y	This field is used to indicate the build in which the fix first appears. In most cases, after the engineer fixes the defect, the field will be set to "Next Build." This field changes to the correct build when that version is actually built.	For this example, leave the field empty.
Responsibility	N	The responsibility field enables the selection of the person who should act on the defect. Depending on the position of the CR in the CR life cycle, this person could be a developer, a QA engineer, or the person who first reported the CR.	For this example, either leave this field blank, or assign it to the lead engineer on the project, who will assign it to the appropriate person.
Addressed by	Y	This field is automatically filled with the name of the person who originally wrote up the CR. It is not editable.	NA
Description/ Steps to Reproduce	Y	Select the Description Tab. In the Description/ Steps to Reproduce field, enter detailed information about the defect. Specifically, the description should build on the synopsis information. The Steps to Reproduce information is the most important data entered in the CR because it provides a detailed step-by-step method of reproducing the defect. The more detailed the information, the more likely the responsible developer will be able to determine the cause of the defect and fix the defect.	Steps to Reproduce might look as follows: 1. RClick the column headers in the upper pane. 2. Select Show Fields. 3. LClick Show Advanced Fields check box. // The check box is activated. 4. LClick Show Advanced Fields check box. // The check box is deactivated. 5. EXP: The standard fields appear in the Available Fields list box. 6. ACT: No fields appear in the Available Fields list box.

Description Syntax/Shorthand

When using a defect tracking system, most organizations employ a special syntax/ shorthand to describe the steps required to reproduce a problem. The following table lists a recommended shorthand and describes each item in that shorthand.

Table 15.4 Change Request Syntax/Shorthand

Shorthand	Description
1, 2, 3,...	Steps are numbered.
LClick	Click the left mouse button.
RClick	Click the right mouse button.
DClick	Double click the left mouse button
[ ]	A keyboard button to be pressed. For example, *[F1*] for Help or [*F5*] for Refresh.
< >	A dialog button. For example: Press <OK> or <Cancel>.

Table 15.4 Change Request Syntax/Shorthand (continued)

Shorthand	Description
>	Menu separator. For example: Select File > Open *or* Topic > Tools > Reports.
//	A comment that is not an actual step. For example: 3) LClick the field. // At this point my machine started to smoke.
EXP	Expected results. For example: EXP: The focus moves to the next field.
ACT	Actual results. For example: ACT: The application crashes.

Appendix

G

Glossary

access rights	A security feature. The rights granted (or denied, in exceptional cases) to users or groups that allow them to see items, modify them, create them, and perform other functions on them.
add files	The process of placing files under version control by adding them to a project view.
administrator	A functional role in the operation of the application. Project administrators create and/or import projects, manage server configurations, manage users and groups, and perform other functions.
All Descendants button	When selected, the view window displays information for the selected folder and its child folders. When not selected, the view window displays only the items associated with the folder and not with its child folders.
alphanumeric	A value consisting of the letters A-Z and the digits 0-9.
alternate working folder	An alternate working folder enables you to store working files on your workstation at a location you specify. Creating an alternate working folder for the root of a view or a branch in a folder hierarchy can alter the paths of the working folders for child folders.
archive	The file or group of files that make it possible to recreate past revisions of a file that is under version control.
ASCII file	A text file. The application recognizes a text file by its lack of null characters.
audit entry	A record of an action performed on a project that appears in the audit log. For example, every time a file is added to a project view, that action is entered in the audit log.
audit log	A chronological record kept by the application showing all actions performed on folders and items.
author	The person who created the revision.
automatic refresh	An automatic update feature of the application. Because projects, folders and items managed by the application are subject to continuous revision by various team members, screen information can become outdated after it is first displayed. The Automatic Refresh function periodically reads the project database to update the displayed information.

binary file	Any file that is not strictly an ASCII text file. Binary files contain programs or data in machine code. Examples of binary files and their respective file name extensions are executables (.exe and .dll), word processing documents (.doc in Word for Windows), spreadsheets (.xls in Excel), object files (.obj) and bitmaps (.bmp).
blank branching view	An empty branching view. This type of view has no correlation to its parent view.
branch	The creation of an independent item that is derived from a corresponding item in a parent view. In the case of a text file, the branched item can later be merged with the file from which it originated. See also *branching view*. Also a branch of a tree, such as the folder hierarchy or a topic tree.
branch revision	The revision number is an identification number assigned by the application to a revision of an item. It indicates how many consecutive revisions have been made in a view, and it branches when the item is branched into a new view. For example, if a branching view includes a change request with the revision number 1.9, the next revision of the change request in the new view will be 1.9.1.0.
branching view	A view that may or may not be derived from a parent view. When not derived from a parent view, it is a blank branching view. Branching views always permit branching. If branching views float and have the Branch On Change option set, they are updated by the parent view on a file-by-file basis until the file changes in the branching view. If they float and do not have the Branch On Change option set, updates are sent to the parent view until the Branch On Change option is set. If these views are based on a label, a promotion state, or a moment in time, they are read-only unless the Branch On Change option is set.
build	The process of compiling, assembling, and linking all project files in the proper sequence to produce the product. Also an event in the life cycle of a product that represents a quantifiable step in progress for a project. For example, a software product may deliver a new build every few days to the team members who test the product. Applying a build label to a set of revisions at the view level ensures that every item in that view at the time of the build has the same label, which makes it possible to identify the source code file revisions that were in each build.
build label	A view label can also be designated as a build label. Build labels change the setting for the Addressed In field for newly resolved change requests from Next Build to the build label. This action lets testers know that a resolved change request can be tested by installing that build. See also *view label.*
change request	Identification of a defect or suggestion for modifying a product.
change request list	After an item has been selected from the folder hierarchy, the list of change displays when you select the Change Request tab. The list may be refined by the Descendants button and the filter you select.
change request number	The number that the application assigns to a change request when it is first submitted.
check-in	The operation that stores a new file revision in the hive archive and cache. The application permits a number of individuals to work on a common set of files by allowing only one team member to revise a project file at a time. A check-in operation marks the end of one revision. The team member who checks in the file can keep it locked or can unlock it and release it, so others can modify it.
check-out	The operation that copies a revision of a file from the project archive (or vault) to a team member's working folder. A user can check out a file with or without the intention of altering that file. The application permits a number of individuals to work on a common set of files by allowing only one team member to revise a project file at a time. Locking the file marks the beginning of one author's revision.

concurrent license	Sometimes called a floating license. Can be used by any user who does not have a named user license. For example, users without named user licenses receive concurrent licenses on a first-come, first-served basis. After all concurrent licenses are in use, users attempting to log on are notified that no more licenses are available at this time.
column header	The label or name for a column that appears in the upper pane of the project view window. Clicking the header initiates sort operations based on the values in that column. Each column header is the name of a property for the items in the list.
command-line file	When you compare two files or revisions, the application sends a command-line file to Visual Diff containing the appropriate options. A similar file is sent to Visual Merge when you merge two files. If there is no space for this file or no temporary folder in which to create it, an error message appears.
component	Major functions of the application are referred to as components. The application includes File, Change Request, Requirement, Task, Topic, and Audit components.
compression	Data that is transferred between your workstation and the server can be compressed. Compression reduces the amount of traffic on the network. However, the time to compress and decompress the data adds to the transfer time.
configuration	A relative arrangement of parts or elements. The application has view, folder, item and server configurations. A view, folder, or item configuration is the isolation of that view, folder, or item to a particular revision based on a point in time. For example, you can configure a view to be: ■ Current - to see the tip revisions of every folder and item in the view. ■ A view label - to see all revisions in the view that have the selected label assigned to them. A view label initially represents a point in time although the label can be adjusted to include revisions that were not current at that point in time and exclude revisions that were. ■ A promotion state - to see all revisions in the view that have been assigned the label that is currently associated with the selected promotion state. ■ A selected date and time - to see all revisions in the view that were current at the specified date and time. You can also configure individual folders and items. See also *server configuration*.
container	A term indicating the ability to contain other types of items. For example, a project is a container for views, folders, and items. Views and folders are also containers.
context menu	A menu that appears when you right-click an area in the project view window. The menu contents vary with the context, that is, the item or area that has been clicked.
Current	File status indicating that the content of the file in the working folder is the same as the content of the tip revision of the file.
defect	A fault or error in a product.
defect tracking	A systematic way of recording information about change requests and maintaining a history of their detection and eventual resolution.
delete	To remove information from the application's databases. You can delete: ■ Items within the project. They can become visible again by rolling back the view or parent folder to a time prior to the deletion. ■ Working files. When these are deleted, they are gone. If the working file was checked out and unchanged, added, and/or checked in before the deletion, no data was lost. See also *remove from version control*.

delta	A progressive revision to a file or a list of changes.
delta storage	A method of computing differences between progressive revisions of a file. In StarTeam 6.0 and earlier, the application could store the first revision of a text file in its entirety in the Archive, along with a set of forward deltas for later revisions. This use of delta storage disappeared with StarTeam 2005. StarTeam 6.0 and earlier releases also used deltas to optimize for slow connections if users selected that option. StarTeam Server 2005 and later releases continue this practice, but store these deltas in each hive's Delta folder, which is a subfolder of each hive's cache path. See also *full revision storage*.
e-mail	An organized system for the delivery of paperless, “electronic” messages. the application provides advantages (such as mailing items) for teams that use an MAPI-compliant e-mail system.
encryption	Data that is transferred between your workstation and the server can be encrypted. Encryption protects files and other project information from being read by unauthorized parties over unsecured network lines.
endpoint	A unique identifier for a server. The endpoint is a port number for the TCP/IP protocol.
exclusive lock	A type of lock that tells others that you are editing a specific file. The icon others see is yellow padlock, while you see an icon with a small yellow key and the head and shoulders of a person. This lock notifies others that you intend to revise the file. Other team members cannot check in a file that you have exclusively locked.
file	A file under application version control. Therefore, a file that is in a project.
file compression	A technique for reducing the size of a file by removing redundant information. Most disk files contain repetitions of common sequences of characters. Compression algorithms remove the additional occurrences of these sequences and save information that permits their restoration. By selecting file compression, you may reduce archive (or vault) space requirements and improve performance.
filter	The criteria used to select a few items from among many. The Filter drop-down list box in the view window enables you to display only the items that are of interest. Applying a filter also controls what columns are displayed, what columns are sorted, and how to group values in the sorted columns.
fixed change request	A designation for a change request indicating that it has been resolved.
float	A view is said to float if it stays current with its parent view; that is, updates from the parent are sent to the child view. If the child is not read-only, updates from the child also go to the parent.
folder	A folder in the application folder hierarchy. Folders, which may have items and child folders, help organize the project view into discrete, understandable parts. For example, a project for a software product might have SourceCode, User Manuals, and Corporate Libraries as folders. Each folder has a working folder on your hard drive that corresponds to it.
folder hierarchy	The hierarchical display of a view and its associated folders. The folder hierarchy is always displayed in the left pane of the view window.
foreign refresh	Pressing *F6* in the Windows client enables you to update the information in your view window for a PVCS or Visual SourceSafe project with which you are communicating via the application. Other clients can access foreign files but cannot refresh them.
frozen	An item is said to be frozen, and therefore, read-only, if it is based on the state of its corresponding item in the parent view at a specific moment in time AND it cannot be branched. An item is also frozen if you reconfigure it to a specific label, promotion state, or time in the past.

full revision storage	A storage method in which all revisions of a file are saved in their entirety. In StarTeam 6.0 and earlier releases, choosing full revision storage reduced the processing time necessary to reconstruct complex binary files stored in delta format in the Archive. In StarTeam 2005 and later releases, all archive files are stored in their entirety, although some may be compressed. See *delta storage*.
history list	The list of revisions for the item selected from the list in the upper pane of the view window. The application displays this list when you select the History tab on the view window lower pane.
Internet Locator Server (ILS) and Internet Locator Service	Microsoft Internet Locator Server provides a dynamic directory service to facilitate real-time communication on the Internet, making it easier for users to find and connect to each other directly using the IP address configured for their PC on the network. It provides a mapping between the static user information and the IP address that is configured for their internet connection. ILS is also a service or server assigned to your IP address to get past the farewell at your location.
item	An application object or element. Items include projects, views, folders, files, change requests, requirements, tasks, topics, and audit entries.
keep current lock status	An option that does not allow the lock status (exclusive, non-exclusive, or unlocked) to be changed as part of the current operation.
keyword	Reserved words beginning and ending with a dollar sign ($). When used in a text file, the application replaces them with the data that they represent. For example, $Author$ is replaced by the name of the user who checked in the file.
keyword expansion	A technique used to insert information in a text file in which keywords are replaced by the data they represent. For example, $Author$ is replaced by the name of the user who checked in the file.
label	The act of attaching a view or revision label for one or more folders and/or items. See build label, revision label, and view label.
labeled configuration	This option limits a new or reconfigured view to items with the view label that you specify when creating the view. If the view is a branching view, the label must be created in this view. If the view is a reference view, it can inherit labels from its parent view. This option is disabled if the view on which the new or reconfigured view is based has no labels defined.
licenses	See named user license and concurrent license.
lock	File locking is a technique to inform others that you are revising a file (exclusive lock) or thinking about revising it (non-exclusive lock). Checking files in and out does not automatically imply locking in the application. You can manually lock or unlock a file during a check-in or check-out operation. However, you can also lock or unlock files with the Lock/Unlock dialog. Locks can be overridden and can be broken. Depending on your personal options, your unlocked files may be marked read-only, which prevents you from making any changes to files that may be locked by others.
MAPI	Acronym for Mail Application Programming Interface. A programming interface that permits an application to send and receive electronic mail via the Microsoft Mail messaging system. The application uses SMTP (Simple Mail Transport Protocol) and not MAPI. See also *SMTP*.
membership	All users are explicit or implicit members of at least one group. Explicit members are those who were specifically assigned to the group. Implicit members are members who belong to the group because they are explicit members of one of the group's child groups.
Merge	The file status that indicates that the working file is not based on the tip revision of the file. When you check this file in or out, the application asks if you want to merge it with the tip revision.

merge	The process of combining a working file with the tip revision of that file or of combining a branched file with the original file from which it was branched. In this three-way comparison, both files are compared with the file revision that is their most recent common ancestor. The combined file can be inspected or revised by the user before it is checked in.
milestone	An event in the life cycle of a product chosen to represent a significant step in progress, for example, the alpha, beta, or final release of a product. In the application, when you reach a milestone, you can apply a view label, usually a build label, to indicate that the milestone has been reached. In Microsoft Project, a milestone is a type of task that represents a significant landmark, development, or turning point in the life of a scheduled project. You usually define a milestone by entering a task name and assigning it a duration of zero. Milestones typically signal that the work has started or has completed; they do not represent the action of doing the work.
Missing	File status indicating that the file is not in your working folder. You may want to check out the file so that its status is updated to Current. See also *Current*.
Modified	File status indicating that the working file has been altered and is based on the tip revision of this file. You may want to check this file in.
named user license	Sometimes called a fixed license. Can be used only by the user who has been assigned that license. For example, if a server has five named user licenses and 25 concurrent licenses, the five users with the named user licenses are guaranteed access to the server. No one else can use their licenses.
non-exclusive lock	A lock indicating that you are thinking about changing the file. It notifies others that they should talk to you before they lock the file exclusively and change it. While you have a file locked non-exclusively, others can check the file in. For text files, it is not unusual for several people to modify the file simultaneously and rely on Visual Merge to manage the changes unless there are conflicts. In that case, a non-exclusive lock helps you identify the people whose input you need to resolve a conflict.
Not In View	File status indicating that the file is in the working folder, but not in the view. You may want to add this file to the view.
notification of events	If you use system tray notification, the application notifies you when an item needs to be read by displaying the appropriate icon in the system tray. If your server is set up for e-mail notification, you receive e-mail notification about items for which you are responsible or for which you are a recipient.
numeric	A value consisting of the valid digits 0-9.
object	A generic term used in object-oriented programming and elsewhere to indicate something upon which operations can be performed. See *item*.
Out Of Date	File status indicating that the working file is a copy of an old revision of the file. If you need the current revision, you should check it out.
owner	A type of information stored by the application for items. On the History tab, it is listed as Author for revisions to items. In addition, it is referenced in the following ways: ■ For a file: The user who added the file to the project view. ■ For a revision: The user who checked in the revision. ■ For a change request: The user who submitted the change request. This is also listed in the Entered by field for a change request. ■ For a topic or response: The user who wrote the topic or response; also listed in the Created by field. ■ For a task: The user who created the task.

pane	The project view window is divided into three panes. The left pane displays the hierarchy of folders in the current view of the project. The right upper pane lists the file, change request, requirement, task, topic, and audit tabs and the right lower pane displays the detail, history, reference, or link information corresponding to the upper right pane.
personal options	User-selectable choices for the behavior of the application on a specific workstation.
profile	For StarTeamMPX: the application has Event Handler profiles that contain options and settings for use with Message Broker and Multicast Services. For a view: A profile is a set of limitations applied to a view. It specifies which folders are visible and which EOL and path case sensitivity settings apply to those folders.
project	A set of related views (each of which usually represents a version of a single product under version control), project properties and project access rights.
promotional model	The application provides an object-oriented architecture that supports entity-relationship modeling. It enables you to move (promote) changes between different stages of the project, for example, from development to testing to product release. Developers work on code changes in promotional states that are isolated from other development efforts.
promotion state	A state through which a product passes. For example, a software application that goes though a development, test, and release cycle could use the promotion states Development, Test, and Release. In the application, each promotion state has a view label assigned to it. The view label can change over time, but users who always work in a specific state may not know the label currently assigned to the state.
promotion state configuration	The basis for a new view or a reconfigured view. The view contains only items with the promotion state you specify. This option is disabled if the view on which the new view or the reconfigured view is based has no promotion states defined.
properties	Attributes stored for each item (and each revision of an item) under version control.
protocol	A set of rules governing how something is done. Network protocols govern how data is transported over the network.
query	The criteria used to select a few items from among many. You can apply a query to items in the upper pane to display only the items that are of interest.
read-only reference view	A reference view that cannot be modified from the view itself, although it may be modified when the parent view is updated. If it floats, it is updated. If it is based on a label (or a promotion state) and the items with the specified label change, the read-only reference view will reflect that. If it is based on a specific date and time, it is frozen as a copy of what the parent view looked like at that point in time.
reference count	A list of the items that reference another item. For example, a file may be shared by two project views on the same server (or even between folders in the same view) and, therefore, have two references to it.
reconfiguring	You can reconfigure a view, file, change request, topic, or task to a point in the past, as defined by a label, promotion state, or a point in time.
reference view	A view derived from a parent view that generally uses a different folder as the root folder and uses the same working folders as the parent. If it floats, it receives updates when the parent view changes. If it floats and is not read-only, it sends updates to the parent view when it changes. If the reference view is based a specific label or date and time within the parent view, it is frozen at that moment in time and is read-only.

refresh	The updating of the information in the project view window or its upper pane. ■ Press *F5* to refresh information about the current view. ■ Press *Shift+F5* to update only the data in the upper pane, such as the files list or change requests list. ■ Press *Ctrl+F5* to update the data and simultaneously collapse all the displayed groups or trees. ■ Press *F6* to refresh the foreign files (files in PVCS or SourceSafe projects).
remove from version control	To make an item invisible and inaccessible in the current view. However, all information about the item remains in the repository and, in the case of files, the hive. You can roll back a folder or view to a time when the item existed if you need to access it again.
repository	Location in which the hive, attachment folder, log files, etc. are stored; associated with a specific server configuration.
repository customization	Modifying values of existing enumerated fields and creating customized fields to represent item information that is specific to your working environment.
Requirement component	The application enables users to create, track, and complete requirements related to the project. This component also imports requirements from CaliberRM.
response	Any of a number of replies to a topic that (along with the topic) form a hierarchical structure called a topic tree.
revision	When an item, such as a change request, is revised, each set of changes is saved as a revision. For example, when a change request changes its status from new to open to resolved, three revisions have occurred.
revision label	A revision label provides a convenient method of identifying a revision of an item or a set of revisions by name. This type of label is s primarily used for files. For example, when you check in a group of files that need to be checked out together, you can give them a revision label.
revision number	The revision number is an identification number assigned by the application to a revision of an item. It indicates how many consecutive revisions have been made after the item was originally created.
root folder	The topmost folder in the folder hierarchy. The root folder in the root view of a project often has the same name as the view and the project.
secondary sort	Sorting items in a list that is already sorted (primary sort). For example a file list might be sorted by extension, then in a secondary sort, sorted by name.
server	A computer or system that provides services to clients, which may be other computers.
server configuration	Contains the repository and option settings selected when the application Server was set up. For example, the administrator may want projects to use encryption, so the server configuration specifies these features. Items, such as folders and files, can be shared if their projects use the same server configuration.
shortcut	A file that starts an application, often with a specific document or set of data. The shortcut is usually stored as an icon on your desktop.
SMTP	Simple Mail Transfer Protocol, commonly used for Internet and UNIX mail systems. It usually uses port 25. Sending items via e-mail in the application requires SMTP.
sort or primary sort	To place items in ascending or descending order in the upper pane of the project view window based on the value in one column. Depending on the values in the column, values are sorted numerically, alphanumerically, or by an internal order or key. Click the column header once to sort. Click a second time to change the sort order from ascending to descending or vice versa. To group sorted items by value, use the *Sort and Group* dialog.

StarDisk	StarDisk is a virtual file system that enables you to use conventional Windows applications, such as Windows Explorer, Microsoft Word for Windows, and Microsoft Developer Studio, to access and manage files that are under version control. You can use StarDisk to map a view to a virtual drive, then access any file on that drive from Explorer or another application. If the file is not checked out, StarDisk can check it out for you.
status	File status is the relationship between the working file and the tip revision in the repository. The file statuses are: In View, Not in View, Missing, Current, Merge, Modified, Out of Date, and Unknown. For a discussion of file status, see the *StarTeam User's Guide.*
Task component	The application enables users to create, track, and resolve tasks related to the project. This component also interoperates with Microsoft Project.
TCP/IP	A protocol for communication between computers used by the Internet. Acronym for Transmission Control Protocol/Internet Protocol.
TCP/IP Sockets	Sockets provide a means or connection for communicating between two different processes. These processes may be on the same machine or they may be on different machines.
team member	Any project users. Sometimes employed to denote a user that does not have administrator status.
test command	A command that you specify so that the application can perform an automated test procedure to determine whether a defect still exists.
text file or ASCII file	A file that contains only printable text characters, spaces, carriage returns, and sometimes tabs and an end-of-file marker, without any formatting codes. The application identifies as text files any files that contain no null characters. All other files are binary.
time stamp	Information maintained by the application about files and revisions. ■ For file revisions: The date and time that the revised file was checked into the application. ■ For files: The date and time that the working file was created.
tip revision	The most recent revision to an item, such as a change request. The tip revision is based on the view configuration. For example, if a view is configured to a particular label, the tip revision is the revision with that label. If a view is configured to a particular date and time, the tip revision is the one that was checked in just prior to that point in time.
topic	The first message on a particular subject attached to a folder in the folder hierarchy. After being submitted by one team member, others may respond to it, which creates a topic tree.
topic tree	A hierarchical structure that appears in the project view window when the Topic tab and the Tree Display button are selected. Its root is a topic and its branches consist of responses to that topic. A topic tree represents a threaded conversation.
Unknown	File status indicating that the file in the working folder has the same name as a file in the view but that it was not checked out of the repository. You may have copied it from another location. Use Update Status to determine the correct status.
unlock	The process of releasing a locked file, indicating that you are no longer changing the file or considering changing it.
user	An individual given access to a server configuration and the projects it manages through the Server. Usually, access is protected by a password. A user is also referred to as a team member.
vault	A folder in which revisions of the files that are under version control are stored. See *archive*.
vault file	A file in the vault that stores file revisions. See *archive*.
variant view	See *branching view*.

version control	Version control is the process of storing and tracking changes (revisions) to one or more files. A version control system maintains the revision history generated as the files evolve into their final forms. The main advantage of an automated version control system is fast, easy recall of previous revisions. The application also tracks revisions of other types of items.
version control system	Application software used to help manage multiple revisions of the same file.
view	Consists of a folder hierarchy, the items associated with each folder in that hierarchy, view properties, and view access rights. The name of a project, its root view, and the root view's root folder are often identical. The behavior of the view may or may not permit branching, may or may not be read-only, and may or may not have a connection to its parent view.
view label	The main purpose of a view label is to time stamp the entire contents of the view. Then you can roll back the view to that label and see everything that existed at that point in time. If the view label is not frozen, you can make changes. You can include or exclude folders and items by attaching or detaching the view label. You can also move a view label from one revision to another.
view window	The window in which the application displays an open view; also called the project view window. This window contains three panes. The left pane displays the folder hierarchy. The upper pane displays the items associated with the folder selected from the folder hierarchy. The lower pane displays details, history, references, and links for the item selected from the upper pane.
Visual Diff	A utility that compares two text files or two revisions of a single text file and shows the differences (if any) between them.
Visual Merge	A utility that helps you merge two revisions of a file that are based on the same checked-in revision. It performs a three-way comparison, comparing the two revisions to each other and to their common ancestor.
work around	Steps required to avoid a defect in a product. Work arounds are frequently used until the defect has been corrected in a new version of the product.
working file	Any file in a working folder can be considered a working file. Often, the file has been checked out for modification. When the file is checked back in, the working file becomes a revision.
working folder	Every application folder has a corresponding working folder to which the working copies of files are checked out and from which files are added and checked in.

Index

Symbols

$basefile
  merge utility parameter 115
$branchtip
  merge utility parameter 115
$file1
  comparison utility parameter 115
$file2
  comparison utility parameter 116
$resultfile
  merge utility parameter 115
$usertip
  merge utility parameter 115
.cdl file contents 282
?
  command-line option 126, 170

A

aborted operations 92
access option
  starteamserver command 257
access rights
  change requests 209
  child folders 206
    attach/adjust revision labels 207
    attach/adjust view labels 207
    change access rights 206
    change behavior or configuration 207
    create and place folder 207
    create links 207
    delete from folder 206
    delete links 207
    detach revision labels 207
    detach view labels 207
    modify links 207
    modify properties 206
    see history 206
    see item and its properties 206
    see links 207
    share/move out of folder 207
  component-level 213
  components
    change access rights 214
    change filter access rights 215
    change query access rights 215
    create public filters 214
    create public queries 214
    delete filters 214, 215
    delete queries 215
    modify filter properties 215
    modify properties 214
    modify query properties 215
    see filter properties 215
    see object and its properties 214
    see query properties 215
  creating project rights 190
  defined 301
  denying 200, 211
  files 207
    add files to a folder 209
    attach revision label 208
    attach view label 208
    break exclusive lock 208
    change behavior or configuration 209
    change item access rights 208
    check in file 209
    check out file 209
    create links 208
    delete from folder 208
    delete links 208
    detach revision label 208
    detach view label 208
    modify links 208
    modify properties 208
    perform maintenance 208
    see history 208
    see item and its properties 208
    see links 208
    set exclusive lock 208
    share/move out of folder 209
  filters 87, 215
  folder level 198
  folders 204
    attach/adjust revision label 205
    attach/adjust view labels 205
    change item access rights 205
    create links 205
    delete from folder 205
    delete links 205
    detach revision labels 205
    detach view labels 205
    modify links 205
    modify properties 205
    see history 205
    see item and its properties 205
    see links 205
  granting 201
  item level 199
  items
    attach/adjust revision labels 210
    attach/adjust view labels 210
    break exclusive lock 210
    change behavior or configuration 211
    change item access rights 210
    create and place in folder 211
    create links 210
    delete from folder 210
    delete links 210
    detach revision labels 210
    detach view labels 210
    modify links 210
    modify properties 210
    perform maintenance 210
    see CR history 210
    see item and its properties 210
    see links 210
    set exclusive lock 210
    share/move out of folder 211
  moving folders and items 199
  project level 196
  projects 82, 202
    delete object 202

promotion states 212
change object access rights 212
modify label assignment 212
queries 87, 215
change object access rights 216
delete object 216
modify object and its properties 216
see object and its properties 216
requirements 209
security guidelines 216
servers 190
add/modify database schema 190
administer components 191
administer user accounts 190
change server access rights 190
change user/operation time 190
create projects 190
modify system configuration 190
modify system policy 190
remotely administer server 190
view server diagnostics 190
view server log 190
view system configuration 190
view system policy 190
setting 201
setting privileges as override 193
sharing folders and items 200
tasks 209
topics 209
understanding 195
using 189
view level 197
views 147
create revision labels 204
create view 204
create view labels 204
define promotion model 204
delete revision labels 204
delete view labels 204
modify revision labels 204
modify view label properties 204
see object and its properties 204
account statuses 65
accounts
deleting users 66
reactivating 67
removing 66
suspending 67
suspending users 67
actions
Add to source 99
Check children 99
Delete from source 98
Delete from target 99
lower pane 97
Manual merge 98
Move from source to target 98
Move from target to source 99
No action 98, 99, 100
Use source file 98, 99
Use target file 98
Add files to the source 113
Add files to the target 111
Add folders to the source 112
Add folders to the target 111
Add projects to the source 112
Add projects to the target 111
Add to source 99
Add views to the source 112
Add views to the target 111
adding
files (defined) 301
folders 154
groups 51, 52
servers 102
users 51, 55
views at command line 133
addresses
e-mail 55
administrators
defined 301
locked out of server 72
Administrators group
defined 51
All Descendants buttons
defined 301
all option
starteamserver command 257
All Users group
defined 51
alphanumeric
defined 301
Alternate comparison utility 115
Alternate merge utility 115
Alternate Property Editors
See APEs
alternate working folders
defined 301
APEs
configuring 81
APPEND 284
Application Log
Windows Event Viewer 19
archive option
starteamserver command 257
archives
defined 301
vault file lock time limit 33
ASCII files
defined 301, 309
assigning
access rights to projects 82
access rights to views 147
view access rights 147
asynchronous
VSS refreshes 226
attach option
starteamserver command 258
attachments
backing up 249
audit log 30
audit entries (defined) 301
defined 301
purging 30
author
defined 301
Automatic merge 98
automatic refresh
defined 301
auto-merge 166
autorecover option
starteamserver command 258

B

b
  command-line option 126, 170
backing up data 249
backups
  need for 90
batch
  command-line option 126
behavior
  shared items 146
bidirectional
  run operation 94
  synchronizations 104
bidirectional synchronizations
  folders 94
  views 95
Binary 118
binary files
  defined 302
  merging 164
  resolving differences 121
blank branching views 140
  defined 302
bold convention 3
brackets convention 3
branch revision
  defined 302
branching
  defined 302
  PVCS support 233
  VSS support 220
branching views 138
  affects on change request workflow 139
  blank 140
  defined 302
  folder-level access rights 199
  replicating 91
Break lock if file needs to be checked in 117
breaklock
  command-line option 126
build labels
  defined 302
build numbers
  viewmerge 178
builds
  defined 302

C

Cache Cleanup Interval 33
cache option
  starteamserver command 258
Catalog Export Utility 276
Catalog_Fields table 276
Catalog_Tables table 276
caveats 90
cdl file contents 282
CFG files 230
change requests
  access rights 209
  affected by change in view 139
  change request list (defined) 302
  change request numbers (defined) 302
  compared in view comparison window 155
  comparing properties in view comparison window 161
  deleting 161
  description syntax/shorthand 299
  entering 297
  merging properties 162, 166
  priority 298
  reconciling 166
  references in view comparison window 168
  reviewing properties in view comparison window 168
  sample guidelines 297
  severity 298
  sharing 160
  status 298
  workflow affected by change in view 139
  workflow after merge 163
  workflow in branching views 139
ChangeRequestsCaching
  starteam-server-configs.xml option 244
changing
  logon passwords 61
  passwords 69
  session options 39
  source view in view comparison window 153
  target view in view comparison window 153
  view's working folders 146
Check children 99
Check-in failed 100
check-ins
  defined 302
  forcing reason 78
  forcing users to lock files first 78
check-out
  defined 302
child folders
  access rights 206
  compared in view comparison window 154
  comparing properties 161
  deleting 161
  merging properties 162, 166
  references in view comparison window 168
  reviewing properties in view comparison window 168
  right to attach revision labels 207
  right to attach view labels 207
  right to change access rights 206
  right to change behavior or configuration 207
  right to create folder and place it in a view or folder 207
  right to create links 207
  right to delete folder from parent folder or view 206
  right to delete links 207
  right to detach revision labels 207
  right to detach view labels 207
  right to modify links 207
  right to modify properties 206
  right to see folder and its properties 206
  right to see history 206
  right to see links 207
  right to share/move out of folder 207
  sharing 160
child views
  included at command line 126
  replicating 92

CipherName
starteam-server-configs.xml option 244
CipherSource
starteam-server-configs.xml option 244
CipherTest
starteam-server-configs.xml option 244
ClientLicenses.st 243
closing
servers 16
cmp
command-line option 126
Collaboration Users group for StarDisk
defined 52
collapsing
folder hierarchy 93, 96
column headers
defined 303
columns
Action 97
Locked By 96
Name 96
Preview Status 97
Result 97
Run Status 97
Source Revision 97
Target Revision 97
Command
starteam-server-configs.xml option 244
command line
merging 169
View Manager 123
command-line
syntax 124
syntax conventions 124
command-line files
defined 303
command-line interface
? option 126, 170
adding views 133
b option 126, 170
batch option 126
breaking locks 126
breaklock option 126
cmp option 126
compressing data 126
configuring source view to date and time 174, 175
configuring source view to promotion state 133, 175
configuring source view to view label 132, 175
conflicts ignored at command line 126
copy option 126
copying history 128
copying revision labels 131
copying times 130
copying user names 131
copyright notices 129
copy-to-target operations 126
cv option 126
deleting folders 171
deleting items 127, 171
df option 171
di option 127, 171
diffopts options 127
diffutil options 127
dry option 171
dry run of viewmerge 171
encrypt option 128, 171
f option 128
help 126, 170
hist option 128
including child views 126
itemtype option 172
limiting merge to one folder 172
lock option 128
locking files 128
locking servers 128
lockservers option 128
mergeopts options 129
mergeutil options 129
no check-in of merged files 172
no Visual Merge 172
noautomerge option 129
noci option 172
noconflicts option 172
nologo option 129
nomerge option 129, 172
norecusion option 172
open option 129
out-of-date files 173
outofdate option 173
p option 129
password files 130, 173
postpone option 130
postponing merges until end 130
promotion states 130
properties file 129, 132
pstates option 130
pwdfile option 173
pwdfile parameter 130
recording merge points 130
recordmp option 130
reptimes option 130
repusers option 131
resume 131
resume option 131
retries option 131
retrying current operation 131
retrywait option 131
rlabels option 131
s option 131, 173
save option 132
sf option 132, 174
sharing folders 132, 174
sharing items 132, 174
si option 132, 174
source option 174
source view 174
specifying comparison utilities 127
specifying folders 128, 132, 133
specifying initial target folder 175
specifying itemtypes to be shared and/or deleted 172
specifying merge utilities 129
specifying project and view 173
specifying projects 129, 133, 134
specifying servers 131
specifying views 133, 134
srcfolder option 132
srclabel option 132, 175
srcproject option 133
srcstate option 133, 175
srctime option 174, 175
srcview option 133

stopping automatic merge 129
stopping merge 129, 172
substitute source files for merged files 134, 176
substitute target files for merged files 135, 176
sv option 133
syntax 169
target option 175
target view 175
targetpath option 175
tgtfolder option 133
tgtproject option 134
tgtview option 134
unlocking servers 134
unlockservers option 134
usemp option 134, 175
usemptgt option 134
usesource option 134, 176
usetarget option 135, 176
using last recorded merge 134, 175
waiting between retries 131
white space 126, 170
command-line utilities
View Manager 89
comparing
change request properties 161
change requests in view comparison window 155
child folders in view comparison window 154
file contents 162
file properties 162
files in view comparison window 154
folder contents 154
folder properties 161
item properties 161
task properties 161
tasks in view comparison window 155
topic properties 161
topics in view comparison window 155
views 150, 152
overview 150
comparison utilities
alternate 114
specified at command line 127
Visual Diff 114
Comparison utility 115
component
defined 303
component-level
access rights 191
components
access rights 213
right to change filter access rights 214, 215
right to change query access rights 215
right to create public filters 214
right to create public queries 214
right to delete filters 214, 215
right to delete queries 215
right to modify filter properties 214, 215
right to modify query properties 215
right to see filter properties 214, 215
right to see query properties 215
compressing
data at command line 126
compression
defined 303
CompressionOffFileExt
starteam-server-configs.xml option 244
CompressionOffFileSize
starteam-server-configs.xml option 245
ComputerName
starteam-server-configs.xml option 245
concurrent license
defined 303
concurrent licenses 59
configuration files 120
configuration options
setting 21
ConfigurationName
starteam-server-configs.xml option 245
configurations
based on promotion states 86
defined 303
PVCS files 230
configuring
APEs 81
source view 102
views at command line 132, 133, 174, 175
conflicts
finding 120
locking files at the command line 172
resolving 121
ConnectionManager.ini 243, 244
containers
defined 303
conventions
bold 3
brackets 3
italics 3
menu selections 3
vertical bars 3
converting
one database to another 274
PVCS files to StarTeam files 239
VSS files to StarTeam files 226, 239
copy
command-line option 126
Copy file history 116
copying
promotion states at command line 130
revision labels at command line 131
server log data 44
times at command line 130
tip revisions 92
user names at command line 131
XML data 281
copyright notices
suppressed at command line 129
copy-to-target
operations 105
run operation 94
copy-to-target operations
at command line 126
folders 95
views 95
CR_MAPPING_ PLACEMENT
defined
282
CR_MAPPING_FIELD 284
defined 282
CR_MAPPING_PLACEMENT 284
CR_MAPPING_STRING 284
defined 282

creating 85
access rights exceptions 211
project with drag-and-drop 76
projects from other applications 234
projects from VSS 221
projects via New Project wizard 75
promotion states 85
property fields 180
read/write reference views 137
read-only reference view 137
reference view 137
views via New View wizard 138
Creation times 109
creation times
problems with 90
credentials 16
Current
defined 303
status 98
customization
defined 308
cv
command-line option 126

D

data
backing up 249
compressing at command line 126
encrypting at command line 128, 171
refreshing PVCS files 238
refreshing VSS files 226
database connections
verifying 41
database information
reviewing 41
databases
backing up 249
exporting catalog tables as comma-delimited files 276
migrating from one to another 274
date/time
creating customization types 182
DBCreated
starteam-server-configs.xml option 245
DBPassword
starteam-server-configs.xml option 245
DBServerName
starteam-server-configs.xml option 245
DBType
starteam-server-configs.xml option 245
DBUserName
starteam-server-configs.xml option 245
default credentials 16
defect tracking
defined 303
defects
defined 303
Delete files from the source 112
Delete files from the target 113
Delete folders from the source 112
Delete folders from the target 113
Delete from source 98
Delete from target 99
deleting
change requests 161
child folders 161
defined 303
files 161
folders 154
groups 54
items 161
itemtypes at command line 172
projects 82
server configurations 20
tasks 161
topics 161
users from groups 64
users from groups temporarily 67
users' accounts 66
views 147
deletions
effect on object IDs 286
effect on user IDs associated with items 287
effect on view configurations 286
effects on reference trees 287
delta storage
defined 304
DeltaStorageOffFileExt
starteam-server-configs.xml option 246
DeltaStorageOffFileSize
starteam-server-configs.xml option 246
demoting
view labels in promotion states 86
denying
access rights 200, 211
descriptions
project 77
projects 77
developer support 3
df
command-line option 171
di
command-line option 127, 171
differences
between views 158
due to deletions 286
due to moved items 287
due to shared items 287
finding 120
resolving 121
diffopts
command-line option 127
diffutil
command-line option 127
disabling
audit log 30
checks for ownership and privileges 193
e-mail 28, 30
server configurations 18
users' accounts 67
displaying
items 156
documentation set 1
dry
command-line option 171
dsn option
starteamserver command 258

E

edit option
  starteamserver command 259
editing
  starteam-client-options.xml 248
e-mail
  defined 304
  enabling 28, 30
  users' addresses 55
empty files
  after merge 178
enabling
  audit log 30
  e-mail 28, 30
  keyword expansion 77
  server configurations 18
  users' accounts 67
encrypt
  command-line option 128, 171
encrypting
  data at command line 128, 171
encryption
  defined 304
endpoints 34
  default TCP/IP 21
  defined 304
enforcing
  process rules 80
entering
  change requests 297
enumerated
  creating types 181
Equivalent Name
  status 99
eval option
  starteamserver command 259
event handlers
  MPX 35
  profiles 35
Event Viewer
  Application Log 19
  used with server configurations 18
exceptions
  denying access rights 211
exclusive locks
  defined 304
Exclusively lock servers on preview 116
Exclusively lock servers on run 116
Exists In Both
  status 99
expanding
  folder hierarchy 93, 96
expiration
  passwords 70
explicit members 53, 63
export utility
  installing 280
  running 281
EXPORT_RECORD_FILE
  defined 282
exporting
  database catalog tables as comma-delimited files 276
  folders 293
  projects 289
  servers 288
  views 291
EXTRACT_DISABLED_PROPERTIES
  defined 283

F

f
  command-line option 128
fax
  telephone numbers 56
file names
  changing 92
files
  access rights 207
  affected by change in view 139
  backing up revisions 249
  breaking locks at command line 126
  comma-delimited 276
  compared in view comparison window 154
  comparing contents in view comparison window 162
  comparing properties in view comparison window 162
  configuration 120
  conflicts ignored at command-line 126
  deleting 161
  deleting at command line 127
  empty after merge 178
  file compression
    defined 304
  finding unresolved files 120
  forcing reason for check-in 78
  forcing users to lock files before check-in 78
  locked by others 104
  locking at command line 128, 172
  log 107
  merging contents 164
  merging properties 164
  options 116
  out-of-date 173
  passwords at command-line 130, 173
  pinned in VSS 222
  read-only 79
  reconciling 166
  references in view comparison window 168
  renaming 92
  replicating 90
  resolving 121
  restart 108
  right to add files to a folder 209
  right to attach revision label 208
  right to attach view label 208
  right to break exclusive lock 208
  right to change behavior or configuration 209
  right to change item access rights 208
  right to check in file 209
  right to check out file 209
  right to create links 208
  right to delete from folder 208
  right to delete links 208
  right to detach revision label 208
  right to detach view label 208
  right to modify links 208
  right to modify properties 208
  right to perform maintenance 208
  right to see history 208

right to see item and its properties 208
right to see links 208
right to set exclusive lock 208
right to share/move out of folder 209
sharing 160
sharing at command line 132
srm 102
starteam-server-configs.xml 244
unlocking 79
working 97
FilesCaching
starteam-server-configs.xml option 246
filters
access rights 87, 215
access rights at component level 213
component-level access rights 191
defined 304
rights at component level 214
finding
differences between views 158
locked files 104
unresolved files 120
fixed change requests
defined 304
fixed license
defined 306
fixed licenses 59
float
defined 304
floating
views 91
floating license
defined 303
floating licenses 59
FOLDER
defined 283
folder
defined 304
folder hierarchy
collapsing 93, 96
defined 10, 304
expanding 93, 96
folder hierarchy, explained 74
folder scope 293
limitations 296
folders
access rights 198, 204
adding to view 153
affected by change in view 139
bidirectional synchronizations 94
compared in view comparison window 154
comparing properties in view comparison window 161
copy-to-target operations 95
deleting 161
deleting at command line 171
deleting from view 154
exporting 293
importing 295
merging properties 162, 166
missing-objects-only synchronizations 95
operations on 94
preview operations 94
properties 154
references 154
references in view comparison window 168
reviewing properties in view comparison window 168
right to attach revision label 205
right to attach view labels 205
right to change access rights 205
right to create links 205
right to delete folder from parent folder or view 205
right to delete links 205
right to detach revision labels 205
right to detach view labels 205
right to modify links 205
right to modify properties 205
right to see folder and its properties 205
right to see history 205
right to see links 205
sharing 160
sharing at command line 132, 174
specified at command line 128, 132, 133
view's working folder 146
VSS refreshes 226
forcing
password changes 69
reason for check-in 78
users to lock files before check-in 78
users to log off 68
foreign archives 90
foreign refresh 226, 238
defined 304
frozen
defined 304
full revision storage
defined 305

G

general
limitations 286
general options 108
getting
ready 90
granting
access rights 201
graphical user interface
View Manager 89
group privileges 193
acquire object ownership 194
change object access rights 194
create object and place it in a container 194
delete object 194
Grant all specific class-level rights for all classes of objects 194
modify object properties 194
purge object 194
see object and its properties 194
groups
adding 51, 52
Administrators 51
All Users 51
Collaboration Users (for StarDisk) 52
deleting 54
deleting users 64
deleting users from all groups 66
deleting users temporarily 67
membership 63
parent 53
properties 195

rearranging the Groups tree 53
Security Administrators 51
System Managers 51
GUI
View Manager 89

H

help
at command line 126
help option
starteamserver command 259
hierarchies
explained 74
hist
command-line option 128
history
copying at command line 128
history list
defined 305

I

icons
property fields 180
view comparison window 155
Ignore whitespace 115
ignoring
ownership 193
privileges 193
implicit members 53, 63
import utility
running 281
IMPORT_RECORD_FILE
defined 283
importing
folders 295
projects 290
servers 289
views 292
inactivity timeout 27, 28
Include child folders 109
Include child views 109
INCLUDE_USER_DATA
defined 283
initialization files 243
Initialized
starteam-server-configs.xml option 246
INPUT_DIRECTORY
defined 283
installing
export utility 280
integer
creating types of 181
Internet Locator Service (ILS)
defined 305
interoperability
StarTeam to VSS 225
VSS to StarTeam 223
introducing
View Manager 89
italics convention 3
item level
access rights 199
item pane
refreshed in view comparison window 168
view comparison 156
items
affected by change in view 139
behavior when shared 146
comparing properties in view comparison window 161
defined 305
deleting 161
deleting at command line 127, 171
displaying 156
dry run of viewmerge at command line 171
mapping numbers from source to target servers 284
merging properties 162, 166
moved 287
reconciling 166
reviewing properties in view comparison window 168
right to attach revision labels 210
right to attach view labels 210
right to break exclusive lock 210
right to change behavior or configuration 211
right to change item access rights 210
right to create and place in folder 211
right to create links 210
right to delete from folder 210
right to delete links 210
right to detach revision labels 210
right to detach view labels 210
right to modify links 210
right to modify properties 210
right to perform maintenance 210
right to see history 210
right to see item and its properties 210
right to see links 210
right to set exclusive lock 210
right to share/move out of folder 211
selected from view comparison window 157
shared 146, 287
sharing 160
sharing at command line 132, 174
itemtype
command-line option 172

K

keep current lock status
defined 305
keyword expansion
defined 77, 305
enabling 77
keywords
defined 305

L

labeled configurations
defined 305
labeling
defined 305
labels
copying at command line 131
defined 305
in branching views 139

replicating 91
view labels in promotion states 85, 86
lengths
passwords 70
licenses
concurrent 59
fixed 59
floating 59
named user 59
licenses option
starteamserver command 259
licensing 57
users 59
limitations
folder scope 296
general 286
project scope 291
server scope 286
view scope 293
limits
logon attempts 71
time of lockouts 71
list option
starteamserver command 259
ListenIP
starteam-server-configs.xml option 246
ListenXML
starteam-server-configs.xml option 246
lock
command-line option 128
Lock unresolved files 118
locked files 104
locking
defined 305
files at command line 128
files with conflicts at the command line 172
forcing user to lock files before check-in 78
servers 103
servers at command line 128
locking options 116
lockout duration 71
locks
breaking at command line 126
lockservers
command-line option 128
log data 44
log files 107
log properties 286
LOG_PROPERTIES
defined 283
logging off
forced by administrator 68
logging on
as different users 16
failure 71
forcing password changes 69
passwords 61
PVCS 235
VSS 223
logging out
See logging off.
logon
attempts 71
sequence timeout 26
statuses 65
logs
purging security event entries 46
View Manager 135
viewmerge 176
lower pane
accessing results 100
actions 97
column headings 96
icons 93
Preview tab 96
properties list 96
results 97
Run tab 96
statuses 97
tree 96
using 96

M

main window
View Manager 92
managing
projects 73
Manual merge 98
Manual merge needed 98
manuals
in documentation set 1
MAPI
defined 305
mapping
item numbers from source to target servers 284
MaxCommandThread
starteam-server-configs.xmloption 246
Maximum cache size 33
members
groups 53
See users.
membership
defined 305
determining 53
users and groups 63
menu selections convention 3
Merge
defined 305
status 98
Merge automatically 117
merge options 113
merge points 175
at command line 134
recording at command line 130
merge utilities
alternate 114
specified at command line 129
Visual Merge 114
Merge utility 115
mergeopts
command-line option 129
mergeutil
command-line option 129
merging
automatic merging stopped at command line 129
binary file contents 164
change request properties 162, 166
child folder properties 162, 166
defined 306
file properties 164

item properties 162, 166
limiting merge to one folder at command line 172
postponed until end of run at the command line 130
recording merge operations 113, 151
specifying initial target folder at command line 175
stopped at command line 129, 172
task properties 162, 166
text file contents 164
topic properties 162, 166
tracking merge operations 113, 151
using last recorded merge at command line 134, 175
views
overview 150
views from command line 169
without check-in at command line 172
without displaying Visual Merge 172
Message Broker 35
migrating
one database to another 274
milestones
defined 306
MinCommandThreads
starteam-server-configs.xml option 247
Missing
defined 306
Missing From Target
status 98
missing object options 110
missing-objects only synchronizations
folders 95
views 95
missing-objects-only
run operation 94
synchronizations 105
MODE
defined 283
Modified
defined 306
status 98
modifying
existing property fields 183
project descriptions 77
project names 77
session options 39
Move from source to target 98
Move from target to source 99
Moved 98, 99
moved
items 287
moving
folders
effects on access rights 199
items
effects on access rights 199
Multicast Service 35

N

name option
starteamserver command 260
named user license
defined 306
named user licenses 59
names
project 77
projects 77
native format 226, 239
new
groups 52
users 55
new option
starteamserver command 260
No action 98, 99, 100
noautomerge
command-line option 129
noci
command-line option 172
noconflicts
command-line option 172
nologo
command-line option 129
nomerge
command-line option 129, 172
non-exclusive lock
defined 306
non-floating views 91
norecursion
command-line option 172
Not In Source
status 99
Not In View
defined 306
notification of events
defined 306
noverifyvault option
starteamserver command 260
Number of retries 110
numbers
mapped from source to target servers 284
numeric
defined 306
numeric types
creating 181

O

object
defined 306
object IDs
different on both servers 286
open
command-line option 129
opening
projects 76
properties files 102
operations
copy-to-target 105
preview 103
resumed at command line 131
resuming 92, 107
retried at command line 131
run 94
unfinished 92
options
Add files to the source 113
Add files to the target 111
Add folders to the source 112
Add folders to the target 111
Add projects to the source 112

Add projects to the target 111
Add views to the source 112
Add views to the target 111
Alternate comparison utility 115
Alternate merge utility 115
Binary 118
Break lock if file needs to be checked in 117
Comparison utility 115
Copy file history 116
Delete files from the source 112
Delete files from the target 113
Delete folders from the source 112
Delete folders from the target 113
Exclusively lock servers on preview 116
Exclusively lock servers on run 116
files 116
general 108
Ignore whitespace 115
Include child folders 109
Include child views 109
Lock unresolved files 118
locking 116
merge 113
Merge automatically 117
Merge utility 115
missing object 110
Number of retries 110
Record each merge 114
Replicate creation times 109
Replicate revision labels 110
Replicate users 109
Replicate view labels 109
replication 109
Resolve differences for binary files 117
Resolve differences for text files 117
setting 108
Text 118
Time between retries ___ minutes 110
Unlock servers after preview 116
Unlock servers after run 116
Use existing file as merged file 117
Use merge tracking for merge resolution 114
Use source as common ancestor 114
Use source file 117, 118
Use specified merge utility 117
Use target as common ancestor 114
Use target file 117, 118
Visual Merge 114
options, server configuration 21, 23
Out Of Date
defined 306
status 99
outofdate
command-line option 173
output
logging 135
OUTPUT_DIRECTORY
defined 284
overview
merging views 150
projects 73
Server Administration 265
View Manager 100
views 137
OVERWRITE 284
owner
defined 306
ownership
ignoring 193

P

p
command-line option 129
p (password) option
starteamserver command 260
pagers
telephone numbers 56
panes
defined 307
refreshing in view comparison window 168
parameters
defined 282
parent groups 53
PASSWORD
defined 284
passwords
at command line 130, 173
expiration 70
forcing changes 69
incorrect logons 71
lengths 70
lockout duration 71
logon 61
PVCS 235
strong 70
time limits 70
VSS 223
personal options
defined 307
phone numbers
See telephone numbers.
PID
starteam-server-configs.xml option 247
planning
server configuration contents 9
pooled connections 31
postpone
command-line option 130
postponing
merges until end of run at the command line 130
PREPEND 284
preview
run operation 94
Preview tab
lower pane 96
previews 103
primary sort
defined 308
privileges 193
acquire object ownership 194
change object access rights 194
create object and place it in a container 194
delete object 194
for groups 193
grant all specific class-level rights for all classes of objects 194
ignoring 193
purge object 194
right to modify object properties 194
right to see object and its properties 194

process rules
enforcing 80
product support 3
profiles
defined 307
event handlers 35
progress bars 106
PROJECT
defined 284
project scope 289
limitations 291
projects 221
access rights 82, 190, 196, 202
adding groups 52
adding users 55
change description 77
change name 77
creating from other applications 234
creating from VSS projects 221
creating via New Project wizard 75
creating with drag-and-drop 76
defined 73, 307
deleting 82
descriptions 77
exporting 289
importing 290
managing 73
names 77
opening 76
overview 73
planning 9
preparations for 73
right to create 190
right to delete project 202
specified at command line 129, 133, 134, 173
VSS refreshes 226
promoting
view labels 85
promotion state configurations
defined 307
promotion states 85
access rights 212
assigning view labels 85
basing a configuration on 86
copying at command line 130
defined 307
demoting view labels 86
example 83
promoting view labels 85
replicating 91
right to change access rights 212
right to modify label assignments 212
setting access rights 212
using 83
promotional models
defined 307
See promotion states.
properties 307
compared in view comparison window 161, 162
folders 154
groups 195
log 286
nodes in lower pane 96
reviewed in view comparison window 168
users 55
properties files
opening 102
saved at command line 132
used at command line 129
property fields
creating 180
creating date/time types 182
creating enumerated types 181
creating numeric types 181
creating text types 182
creating user ID types 183
icons 180
modifying existing properties 183
protocols
defined 307
TCP/IP 35
pstates
command-line option 130
purging
audit log entries 30
security event log 46
PVCS 90
branching support 233
configuration files 230
features not reproduced in StarTeam 236
features reproduced in StarTeam 236
passwords 235
permanent conversion of files to StarTeam 239
preparing for wrapping process 232
refreshing files 238
StarTeam operations not performed in archives 237
StarTeam operations reproduced in PVCS archives 237
StarTeam project 234
wrapped by StarTeam 234
pwdfile
command-line option 173
command-line parameter 130

Q

q (quiet) option
starteamserver command 260
queries
access rights 87, 215
access rights at component level 213
component-level access rights 191
defined 307
right to change access rights 216
right to delete queries 216
right to modify properties 216
right to see properties 216
rights at component level 214

R

r (repository path) option
starteamserver command 260
RC2_CBC 128, 171
RC2_CF 128, 171
RC2_ECB 128, 171
RC4 128, 171
reactivating
users' accounts 67
readiness 90

read-only
files 79
read-only reference views (defined) 307
real property types
creating 181
reason for check-in
forced 78
reconciling
folders and items 166
reconfiguring
defined 307
views in view comparison window 153
Record each merge 114
recording
merge operations 113, 151
recordmp
command-line option 130
reference counts
defined 307
reference views 140
creating read/write 137
defined 307
folder-level access rights 199
replicating 90
references
different on both servers 287
folders 154
in view comparison window 168
refresh
defined 308
refreshes
asynchronous (VSS) 226
automatic, defined 301
per folder 226
per project 226
synchronous (VSS) 226
refreshing 97
data 226, 238
panes in view comparison window 168
PVCS files 238
views 226, 238
refreshing, VSS files 226
reloading
security event log 45
server log 44
remove option
starteamserver command 261
Remove Service button
Server Administration 19
removing
groups 54
users 55
users from groups 64
users from groups temporarily 67
users' accounts 66
removing from version control
defined 308
renaming
files 92
Replicate creation times 109
Replicate revision labels 110
Replicate users 109
Replicate view labels 109
replicating
branching views 91
child views 92
files only 90
labels 91
minimal amount of data 90
promotion states 91
PVCS archives 90
reference views 90
revision labels at command line 131
times at command line 130
user names at command line 131
view labels 91
VSS archives 90
replication options 109
repository
defined 308
verification 272
repository customization
creating property fields 180
defined 308
RepositoryPath
starteam-server-configs.xml option 247
reptimes
command-line option 130
repusers
command-line option 131
REQ_MAPPING_ PLACEMENT
defined 284
REQ_MAPPING_ STRING
defined 284
REQ_MAPPING_FIELD 284
defined 284
REQ_MAPPING_PLACEMENT 284
REQ_MAPPING_STRING 284
Requirement component
defined 308
requirements
access rights 209
requiring
process rules 80
Resolve differences for binary files 117
Resolve differences for text files 117
resolving
conflicts 121
differences between binary files 121
differences between text files 121
Response
starteam-server-configs.xmloption 247
responses
defined 308
restart files 108
restart option
starteamserver command 261
results
Automatic merge 98
Check-in failed 100
lower pane 97
Manual merge needed 98
Moved 98, 99
Source file checked in to target 98, 99
Target file checked in to source 98
resume
command-line option 131
resuming
operation at command line 131
operations 92, 107
retries 110
command-line option 131

retrying
operations at command line 131
retrywait
command-line option 131
reviewing
database information 41
folder properties 154
folder references 154
properties in view comparison window 168
references in view comparison window 168
revision
defined 308
verifying 272
Revision labels 110
revision labels
copying at command line 131
defined 308
revision numbers
defined 308
rlabels
command-line option 131
rolling back
views at command line 132, 133, 174, 175
views in view comparison window 153
root folders
defined 308
rules
process enforcement 80
run operations 94
bidirectional 94
copy-to-target 94
missing-objects-only 94
preview 94
source-to-target 94
Run tab
lower pane 96

S

s
command-line option 131, 173
save
command-line option 132
scope 108
scoped views 137
secondary sort
defined 308
security
privileges 193
PVCS 235
viewing security event logs 45
views 147
VSS 223
Security Administrators group
defined 51
Security Event Log
purging 46
viewing 45
security events
purging from log 46
security guidelines 216
selecting
items from view comparison window 157
serial option
starteamserver command 261
server
adding 13
configuring 13
Server Administration
overview 11, 265
Remove Service button 19
Set Service button 19
Server Administration utility
described 11
server configuration options 21, 23
overriding at the command prompt 256
server configurations
adding groups 52
adding to the Server Administration dialog 21
adding users 55
configuration options (defined) 12
configuration options, setting 21
copying log data 44
defined 308
deleting 20
deleting groups 54
deleting users from all groups 66
deleting users from groups 64
disabling 18
enabling 18
execution mode 11
forcing users to log off 68
locking 17
lockout work around 72
migrating from one database to another 274
overriding sessions options 15
planning 9
printing log data 44
running as a Windows service 18
Service status 19
session options (defined) 12
shutting down 16
starting from Server Administration 15
suspending users' accounts 67
unlocking 17
verifying tip revisions 272
viewing server log 43
server log 44
copying data from 44
viewing the contents 44
viewing the location 44
server scope 288
limitations 286
SERVER_NAME
defined 284
SERVER_PORT
defined 284
ServerGUID
starteam-server-configs.xml option 247
servers
access rights 190
adding 102
adding groups 52
adding users 55
changing session options 39
copying log data 44
copy-to-target operations 95
defined 308
deleting configurations 20
exporting 288
forcing users to log off 68

importing 289
locking 103
locking at command line 128
pooled connections configurations 31
printing log data 44
right to add/modify database schema 190
right to administer components 191
right to administer user accounts 190
right to change server access rights 190
right to create projects 190
right to modify system configuration 190
right to modify system policy 190
right to remotely administer server 190
right to view server diagnostics 190
right to view server log 190
right to view system configuration 190
right to view system policy 190
shutting down 16
special clients access right 190
specified at command line 131
unlocking at command line 134
Service status
server configuration property 19
ServiceMode
starteam-server-configs.xmloption 247
services
Windows, properties 18
session options
changing 39
Set Service button
Server Administration 19
setting 193
access rights 201
inactivity timeout 27, 28
lockout duration 71
logon sequence timeout 26
options 108
privileges 193
vault file lock time limit 33
sf
command-line option 132, 174
shared
items 287
sharing
change requests 160
child folders 160
files 160
files at command line 132
folders
effects on access rights 200
folders at command line 132
items 160
effects on access rights 200
items at command line 132
itemtypes at command line 172
tasks 160
topics 160
shortcuts
defined 308
shorthand
change requests 299
shutting down
servers 16
si
command-line option 132, 174
SMTP
defined 308
sorting
defined 308
source
command-line option 174
Source file checked in to target 98, 99
source files
option at command line 176
parameter at command line 134
source objects
properties files 102
selecting from menu bar 101
selecting from upper pane 93
source view
specified at command line 174
source-to-target
run operation 94
synchronizations 104
special clients
server access right 190
square brackets convention 3
srcfolder
command-line option 132
srclabel
command-line option 132, 175
srcproject
command-line option 133
srcstate
command-line option 133, 175
srctime
command-line option 174, 175
srcview
command-line option 133
srm files 102
standard error 135, 176
standard output 135, 176
starbeamserver options
-cache 258
StarDisk
Collaboration Users groups 52
defined 309
start option
starteamserver command 261
StarTeam
interoperating with VSS 219
See also Cross-Platform Client, Application Server, and Windows Client.
StarTeam Log File 47
customizing 48
data 48
detailed reports 48
events 48
log errors 48
MPX Server events 49
operations 48
purpose 48
StarTeam storage type 226, 239
StarTeam tables
Catalog_Fields 276
Catalog_Tables 276
starteam-client-options.xml 243
editing 248
understanding 248
StarTeamMPX 35

starteamserver
syntax 253
starteamserver options
-access 257
-all 257
-archive 257
-attach 258
-autorecover 258
-dsn 258
-edit 259
-eval 259
-help 259
-licenses 259
-list 259
-name 260
-new 260
-noverifyvault 260
-p (password) 260
-q (quiet) 260
-r (repository path) 260
-remove 261
-restart 261
-serial 261
-start 261
-stop 261
-stoponerors 262
-t (database type) 262
-tcpip 262
-u (user name) 263
-vault 263
-verifyvault 263
-version 263
-view 264
starteam-server-configs.xml 243
contents 244
location 244
starteam-server-configs.xml options
ChangeRequestsCaching 244
CipherName 244
CipherSource 244
CipherTest 244
Command 244
CompressionOffFileExt 244
CompressionOffFileSize 245
ComputerName 245
ConfigurationName 245
DBCreated 245
DBPassword 245
DBServerName 245
DBType 245
DBUserName 245
DeltaStorageOffFileExt 246
DeltaStorageOffFileSize 246
FilesCaching 246
Initialized 246
ListenIP 246
ListenXML 246
MaxCommandThread 246
MinCommandThreads 247
PID 247
RepositoryPath 247
Response 247
ServerGUID 247
ServiceMode 247
Status 247
TasksCaching 247
TopicsCaching 247
UserName 247
starteam-servers.xml 243
starteam-servers.xml file
defined 12
location 12
states
replicating 91
See promotion states.
Status
starteam-server-configs.xml option 247
statuses
accounts 65
Current 98
defined 309
Equivalent Name 99
Exists In Both 99
logon 65
lower pane 97
Merge 98
Missing From Target 98
Modified 98
Not In Source 99
Out Of Date 99
Unknown 100
steps to reproduce 299
stop option
starteamserver command 261
stoponerrors option
starteamserver command 262
strong passwords 70
support 3
suspending
accounts 67
sv
command-line option 133
synchronizations 94
bidirectional 104
missing-objects-only 105
source-to-target 104
synchronous
VSS refreshes 226
syntax
command-line 124, 169
syntax conventions
command-line 124
mutually-exclusive choices 3
options 3
variables 3
System Managers group
defined 51
system tray
e-mail notification 30

T

t (database type) option
starteamserver command 262
tables
database 276
exporting database catalogs 276
target
command-line option 175
Target file checked in to source 98

target files
option at command line 176
parameter at command line 135
target objects
properties files 102
selecting from menu bar 101
selecting from upper pane 93
target view
specified at command line 175
targetpath
command-line option 175
Task component
defined 309
TASK_MAPPING_ FIELD
defined 284
TASK_MAPPING_ PLACEMENT
defined 284
TASK_MAPPING_ STRING
defined 284
TASK_MAPPING_FIELD 284
TASK_MAPPING_PLACEMENT 284
TASK_MAPPING_STRING 284
tasks
access rights 209
affected by change in view 139
compared in view comparison window 155
comparing properties in view comparison window 161
deleting 161
merging properties 162, 166
reconciling 166
references in view comparison window 168
reviewing properties in view comparison window 168
sharing 160
TasksCaching
starteam-server-configs.xml option 247
TCP/IP
default endpoint 21
defined 309
setting minimum encryption for transferred data 35
tcpip option
starteamserver command 262
team members
defined 309
See users.
technical support 3
telephone numbers
fax 56
pagers 56
users 56
voice mail 56
test commands
defined 309
Text 118
text files
defined 309
merging contents 164
resolving differences 121
text types
creating 182
tgtfolder
command-line option 133
tgtproject
command-line option 134
tgtview
command-line option 134
Time between retries ___ minutes 110
time limits
lockouts 71
passwords 70
time stamps
defined 309
time types
creating 182
timeouts
inactivity 27, 28
logon failure 26
logon failures 71
vault file lock 33
times
copying at command line 130
problems with 90
tip revisions
copying 92
defined 309
verifying 272
topic trees
defined 309
TOPIC_MAPPING_FIELD 284
TOPIC_MAPPING_PLACEMENT 284
TOPIC_MAPPING_STRING 284
topics
access rights 209
affect by change in view 139
affected by change in view 139
compared in view comparison window 155
comparing properties in view comparison window 161
defined 309
deleting 161
merging properties 162, 166
reconciling 166
references in view comparison window 168
reviewing properties in view comparison window 168
sharing 160
TopicsCaching
starteam-server-configs.xml option 247
tracking
merge operations 113, 151
troubleshooting 288
turning off
checks for ownership and privileges 193
types
run operations 94

U

u (user name) option
starteamserver command 263
understanding
access rights 195
Unknown
defined 309
status 100
Unlock servers after preview 116
Unlock servers after run 116

unlocking
defined 309
files 79
servers at command line 134
unlockservers
command-line option 134
unresolved files 120
upper panes 93
Use existing file as merged file 117
Use merge tracking for merge resolution 114
Use source as common ancestor 114
Use source file 98, 99, 117, 118
Use specified merge utility 117
Use target as common ancestor 114
Use target file 98, 117, 118
usemp
command-line option 134, 175
usemptgt
command-line option 134
user ID property fields
creating types 183
user IDs 287
user names
copying at command line 131
USERNAME
defined 284
UserName
starteam-server-configs.xml option 247
Users 109
users
account status 65
adding 51, 55
defined 309
deleting accounts 66
deleting from all groups 66
deleting from groups 64
deleting from server configuration 66
e-mail addresses 55
fax numbers 56
group memberships 63
licensing 59
logging on 16
logon status 65
pager numbers 56
properties 55
reactivating accounts 67
removing 55
removing accounts 66
suspending accounts of 67
telephone numbers 56
voice mail numbers 56
usesource
command-line option 134, 176
usetarget
command-line option 135, 176
using
access rights 189
utilities
comparison 114
merge 114

V

variant views
See branching views. 309
vault
defined 309
file lock time limit 33
vault files (defined) 309
verification 272
vault option
starteamserver command 263
vcs.cfg 230
verification
server configuration's vault and database 272
verifying
database connections 41
server configuration's vault and database 272
verifyvault option
starteamserver command 263
version control
defined 310
version control systems (defined) 310
version number
viewmerge 178
version option
starteamserver command 263
vertical bars convention 3
VIEW
defined 284
view comparison
change requests 155
child folders 154
files 154
icons 155
item pane 156
tasks 155
topics and responses 155
view comparison window 152
finding differences 158
refreshing panes 168
selecting items 157
view configurations
deleted objects missing 286
View labels 109
view labels
assigning to promotion states 85
defined 310
demoting 86
promoting from one promotion state to another 85
replicating 91
View Manager
command line 123
logging output 135
view option
starteamserver command 264
view scope 291
limitations 293
view window
defined 310
viewing
security event log activity 45
server log activity 43
viewmerge 169
build number 178
logging output 176
version number 178
views
access rights 147, 197
adding at command line 133
adding folders 153

bidirectional synchronizations 95
blank branching 140
branching 91, 138, 140
changing views in view comparison window 153
child 92
child views included at command line 126
comparing 150, 152
configured at command line 132, 133, 174, 175
configuring 102
copy-to-target operations 95
creating 137
creating read/write reference view 137
creating via New View wizard 138
defined 310
deleting 147
deleting folders 154
exporting 291
importing 292
merged from command line 169
merging 150
missing-objects-only synchronizations 95
operations on 95
overview 137
planning 9
preview operations 95
reference 90, 140
refreshing 226
VSS files 226
refreshing PVCS files 238
right to create revision labels 204
right to create view labels 204
right to create views 204
right to define promotion model 204
right to delete revision labels 204
right to delete view labels 204
right to modify revision labels 204
right to modify view label properties 204
right to see view and its properties 204
rolling back in view comparison window 153
scoped 137
security 147
specified at command line 133, 134, 173
working folder properties 146
zoomed view 137
views, additional 144
views, creating 144
Visual Diff
defined 310
Visual Merge 114
defined 310
voice mail
telephone numbers 56
VSS 90
branching support 220
features reproduced in StarTeam 223
interoperating with StarTeam 219
passwords 223
permanent conversion of files to StarTeam 226, 239
pinned files 222
preparing for wrapping process 219
refreshing files 226
refreshing files in corresponding StarTeam project 226
StarTeam features reproduced in VSS 225
StarTeam projects 221
wrapped by StarTeam 221

W

waiting between waiting
between retries at command line 131
white space
option at command line 126, 170
Windows
Event Viewer 19
windows
view comparison 152
Windows services
properties 18
wizards
New Project 75
New View 138
work arounds
defined 310
workflow
affected by change in view 139
in change requests in branching views 139
in merged change requests 163
working files 97
defined 310
working folders
defined 10, 310
view's working folder 146
wrapping
PVCS 234
VSS 221

X

XML data
copying to target server 281

Z

zoomed views 137